Commodity Trade and the Carry Trade: A Tale of Two Countries*

Robert Ready[†], Nikolai Roussanov[‡] and Colin Ward[§]

October 9, 2013

Abstract

Persistent differences in interest rates across countries account for much of the profitability of currency carry trade strategies. The high-interest rate "investment" currencies tend to be "commodity currencies," while low interest rate "funding" currencies tend to belong to countries that export finished goods and import most of their commodities. We develop a general equilibrium model of commodity trade and currency pricing that generates this pattern via frictions in the shipping sector. The model predicts that commodity-producing countries are insulated from global productivity shocks by the limited shipping capacity, which forces the final goods producers to absorb the shocks. As a result, a commodity currency is risky as it tends to depreciate in bad times, yet has higher interest rates on average due to lower precautionary demand, compared to the final good producer. The model's predictions are strongly supported in the data. The commodity-currency carry trade explains a substantial portion of the carry-trade risk premia, and all of their pro-cyclical predictability with commodity prices and shipping costs, as predicted by the model.

Keywords: carry trade, currency risk premia, exchange rates, international risk sharing, commodity trade
JEL codes: G15, G12, F31

*We benefited from comments by Andy Abel, Rui Albuquerque, Dave Backus, Gurdip Bakshi, John Campbell, Mike Chernov, Ric Colacito, Max Croce, Darrell Duffie, Xavier Gabaix, Jeremy Graveline, Robin Greenwood, Tarek Hassan, Burton Hollifield, Urban Jermann, Karen Lewis, Debbie Lucas, Hanno Lustig, Don Keim, Brent Neiman, Jose Scheinkman, Ivan Shaliastovich, Rob Stambaugh, Andreas Stathopoulos, Sheridan Titman, Adrien Verdelhan, Jessica Wachter, Amir Yaron, Stan Zin, and audiences at the CEPR ESSFM Gerzensee, Minnesota Asset Pricing conference, Oxford-MAN Currency Trading conference, NBER SI, NBIM, SECOR, Texas Finance Festival, SED, WFA, and Wharton. Roussanov acknowledges financial support from the Iwanowski Family Research Fellowship and Wharton Global Research Initiative.

[†]Simon School of Business, University of Rochester

[‡]The Wharton School, University of Pennsylvania, and NBER

[§]The Wharton School, University of Pennsylvania

1 Introduction

A currency carry trade is a strategy that goes long high interest rate currencies and short low interest rate currencies. A typical carry trade involves buying the Australian dollar, which for much of the last three decades earned a high interest rate, and funding the position with borrowing in the Japanese yen, thus paying an extremely low rate on the short leg. Such a strategy earns positive expected returns on average, and despite substantial volatility and a risk of large losses, such as ones incurred during the global financial crisis, exhibits high Sharpe ratios. In the absence of arbitrage this implies that the marginal utility of an investor whose consumption basket is denominated in yen is more volatile than that of an Australian consumer. Are there fundamental economic differences between countries that could give rise to such a heterogeneity in risk?

One source of differences across countries is the composition of their trade. Countries that specialize in exporting basic commodities, such as Australia or New Zealand, tend to have high interest rates. Conversely, countries that import most of the basic input goods and export finished consumption goods, such as Japan or Switzerland, have low interest rates on average. These differences in interest rates do not translate into the depreciation of "commodity currencies" on average; rather, they constitute positive average returns, giving rise to a carry trade-type strategy. In this paper we develop a theoretical model of this phenomenon, document that this empirical pattern is systematic and robust over the recent time period, and provide additional evidence in support of the model's predictions for the dynamics of carry trade strategies.

The fact that carry trade strategies typically earn positive average returns is a manifestation of the failure of the Uncovered Interest Parity (UIP) hypothesis, which is one of the major longstanding puzzles in international finance. It is commonly recognized that time-varying risk premia are a major driver of carry trade profits. In fact, a longstanding consensus in the international finance literature attributed all of the carry trade average returns to *conditional* risk premia, with no evidence of non-zero *unconditional* risk premia on individual currencies throughout most of the twentieth century (e.g. see Lewis (1995)). Consequently, much of the literature has focused on explaining the conditional currency risk premia by ruling out asymmetries (e.g., Verdelhan (2010), Bansal and Shaliastovich (2012), Colacito and Croce (2012)). However, Lustig, Roussanov, and Verdelhan (2011) show that unconditional currency risk premia are in fact substantial; indeed, they account for between a third and a half of the profitability of carry trade strategies.[1] Lustig, Roussanov, and Verdelhan (2011) argue that these returns are compensation for global risk, and the presence of unconditional risk premia implies that there is persistent heterogeneity across countries' exposures to common shocks. In this paper we uncover a potential source of such heterogeneity.[2]

We show that the differences in average interest rates and risk exposures between countries that are net importers of basic commodities and commodity-exporting countries can be explained by appealing to a natural economic mechanism: trade costs.[3] We model trade costs by considering a simple model of the shipping industry. At any time the cost of transporting a unit of good from one country to the other depends on the aggregate shipping capacity available. While the capacity of the shipping sector adjusts over time to match the demand for transporting goods between countries, it does so slowly, due to gestation lags in the shipbuilding industry. In order to capture this intuition we assume marginal costs of shipping an extra unit of good is increasing - i.e., trade costs in our model are convex. Convex shipping costs imply that the sensitivity of the commodity country to world productivity shocks is lower than that of the country that specializes in producing the final consumption good, simply because it is costlier to deliver an extra unit of the consumption good to the commodity country in good times, but cheaper in bad times. Therefore, under complete financial markets, the commodity country's consumption is smoother than it would be in the

[1]See also Bakshi, Carr, and Wu (2008), Campbell, Medeiros, and Viceira (2010), Koijen, Pedersen, Moskowitz, and Vrugt (2012), and Lustig, Roussanov, and Verdelhan (2013) for additional empirical evidence. Theoretical models of Hassan (2010) and Martin (2011) relate currency risk premia to country size. Stathopoulos, Vedolin, and Mueller (2012) assume an exogenous source of heterogeneity in a multi-country model with habit formation.

[2]A number of patterns of heterogeneous risk exposures have been documented empirically. In a pioneering study, Lustig and Verdelhan (2007) show that carry trade risk premia line up with loadings on the U.S. aggregate consumption growth; Lustig, Roussanov, and Verdelhan (2011) and Menkhoff, Sarno, Schmeling, and Schrimpf (2012) link these risk premia to covariances with the global stock market and foreign exchange rate volatility shocks, respectively, while Lettau, Maggiori, and Weber (2013) show that high average return strategies in currency and commodity (as well as equity) markets perform particularly poorly during large U.S. stock market declines.

[3]Trade costs have a long tradition in international finance: e.g., Dumas (1992), Hollifield and Uppal (1997). Obstfeld and Rogoff (2001) argue that trade costs hold the key to resolving several major puzzles in international economics.

absence of trade frictions, and, conversely, the commodity importer's consumption is riskier. Since the commodity country faces less consumption risk, it has a lower precautionary saving demand and, consequently, a higher interest rate on average, compared to the country producing manufactured goods. Since the commodity currency is risky - it depreciates in bad times - it commands a risk premium. Therefore, the interest rate differential is not offset on average by exchange rate movements, giving rise to a carry trade.

We show empirically that sorting currencies into portfolios based on net exports of finished (manufactured) goods or basic commodities generates a substantial spread in average excess returns, which subsumes the unconditional (but not conditional) carry trade documented by Lustig, Roussanov, and Verdelhan (2011). Further, we show that aggregate consumption of commodity countries is less risky than that of finished goods producers, as our model predicts.

The model makes a number of additional predictions that are consistent with salient features of the data. Commodity-currency carry trade returns are positively correlated with commodity price changes, both in the model and in the data (we provide evidence using an aggregate commodity index, which complements the result obtained by Ferraro, Rossi, and Rogoff (2011) who use individual currency and commodity price data). Moreover, the model predicts that conditional expected returns on the commodity-currency carry trade are especially high when global goods markets are most segmented, i.e. when trade costs are particularly high. We show that a popular measure of shipping costs known as the Baltic Dry Index (BDI) forecasts unconditional carry trade returns (but not their conditional component). Our model also rationalizes the evidence of carry trade predictability with a commodity price index documented by Bakshi and Panayotov (2012), since commodity prices are typically high in the model during booms, when trade costs are also high.

In order to evaluate the model's ability to generate quantitatively reasonable magnitudes of currency risk premia and interest rates we calibrate it by allowing for the possibility of very large jumps in productivity - i.e., rare disasters, as in the literature on the equity premium puzzle (e.g., Longstaff and Piazzesi (2004), Barro (2006), Gabaix (2012), Wachter (2013)). The calibrated model is able to account for the observed interest rate differentials and average returns on the commodity currency carry trade strategies without overstating consumption growth volatility or implying an unreasonably high probability of a major disaster.

2 Model

2.1 Setup

There are two countries, each populated by a representative consumer endowed with CRRA preferences over the same consumption good, with identical coefficients of relative risk aversion γ and rates of time preference ρ. The countries differ in their production technologies, each specializing in the production of a single good. The "commodity" country produces a basic input good using a simple production technology

$$y_c = z_c l_c^{\alpha};$$

assuming one unit of commodity country's non-traded input l_c (e.g., labor, land, etc.) is supplied inelastically, so that this is equivalent to an exogenous endowment of basic commodity equal to the productivity shock z_c ($y_c = z_c$).

The "producer" country only produces a final consumption good using basic commodity input b and labor:

$$y_p = z_p b^{1-\beta} l_p^{\beta},$$

which is subject to a productivity shock z_p, with one unit of producer country's non-traded input also supplied inelastically.

The countries are spatially separated so that transporting goods from one country to the other incurs shipping costs. Our model of shipping costs extends the variable iceberg cost of Backus, Kehoe, and Kydland (1992), where each unit of good shipped in either direction loses a fraction

$$\tau_i(x, z_k) = \kappa_0^i + \kappa_1^i \frac{x}{z_k},$$

which depends on the total amount of goods shipped in the same direction, x, and the shipping capacity available at time t, z_k. For simplicity we assume that this shipping capacity (or, equivalently, shipping sector productivity) is exogenous (although a model with investment in shipping capacity yields similar implications). Since the costs of shipping raw commodities and manufactured goods are likely to be different, we allow two sets of parameters ($i \in c, f$).

Since the commodity country has no alternative use for the basic good it produces, in equilibrium all of its supply is shipped to the producer country. Total output of the final consumption good is therefore

$$y_p = z_p[z_c(1 - \tau_c(z_c, z_k))]^{1-\beta} l_p^\beta.$$

In the producer country, the representative competitive firm solves

$$\max_{l_p \in [0,1]} \pi_p = z_p(z_c(1 - \tau_c(z_c, z_k)))^{1-\beta} l_p^\beta - w_p l_p - P z_c(1 - \tau_c(z_c, z_k)),$$

where w_p is the wage paid to labor and P is the price of one unit of basic commodity. From the first-order conditions and zero profits, the price of the basic commodity is given by

$$P = \frac{(1-\beta) y_p}{(1 - \tau_c(z_c, z_k)) y_c} = (1-\beta) z_p [z_c(1 - \tau_c(z_c, z_k))]^{-\beta}.$$

Consumption allocations for the commodity country and the producer country, c_c and c_p, are determined by the output of the producer country y_p and the amount X of final consumption good exported to the commodity country. We will consider complete financial markets as our benchmark case, so that equilibrium consumption allocations to the two countries over time and across states of nature will be determined as a result of a risk-sharing arrangement, and the real exchange rate is pinned down by the absence of arbitrage in the financial markets (as well as the markets for the consumption good). In contrast, in (financial) autarky, whereby trade is balanced in every period since trade in financial claims is impossible, the producer country consumption equals to its share of output equal to $\beta z_p[z_c(1 - \tau_c(z_c, z_k))]^{1-\beta}$ (if labor is the only non-traded factor, this quantity represents the total wage bill in the competitive equilibrium), while the remainder of the output is exported to the commodity country in the form of payment for the basic commodity ($X_{aut} = (1-\beta) z_p [z_c(1 - \tau_c(z_c, z_k))]^{1-\beta}$), which implies that after trade cost the commodity country income/consumption would equal $X_{aut}(1-\tau_f(X_{aut}, z_k))$. The real exchange rate in this case is determined by the terms of trade (i.e., the relative price of the basic commodity).

The production economy outlined here is very simple (e.g., it is essentially static, as there are no capital or other inter-temporal investment margins), intended to highlight the main mechanism based on the interplay of specialization and trade costs. Gourio, Siemer, and Verdelhan (2013) and Colacito, Croce, Ho, and Howard (2013) study currency risk premia in fully dynamic production economies that could potentially be generalized to incorporate the type of heterogeneity we consider.

2.2 Dynamics

We assume that the shocks to productivity experienced by the final good producer are permanent, so that its evolution (in logs) follows a jump-diffusion process:

$$d\log z_{pt} = (\mu - \mu_Z \eta)\, dt + \sigma_p dB_{pt} + dQ_t.$$

Let $N(t)$ be a Poisson process with intensity η, and let $-Z_1, -Z_2, \ldots$ be a sequence of identically distributed random variables drawn from a truncated Pareto distribution with minimum jump Z_{min}, maximum jump Z_{max}, and shape parameter α. Denote this distribution's mean as μ_Z. Define the compound Poisson process:

$$Q(t) = \sum_{j=1}^{N(t)} Z_j = \int_0^t Z_s dN_s,\ t \geq 0.$$

$$\Rightarrow dQ(t) = Z_{N(t)} dN_t,$$

so that μ is the uncompensated drift of the jump-diffusion, and the growth rate of the productivity shock process can be written as

$$\begin{aligned}\frac{dz_{pt}}{z_{pt^-}} &= \left(\mu - \mu_Z \eta + \frac{1}{2}\sigma_p^2\right) dt + \sigma_p dB_{pt} + (e^{Z_{N(t)}} - 1) dN_t \\ &\mathrel{\mathring{=}} \mu_p dt + \sigma_p dB_{pt} + (e^{Z_{N(t)}} - 1) dN_t,\end{aligned}$$

where $z_{pt^-} = \lim_{s\uparrow t} z_{ps}$ is the process's left-limit, a convention used throughout.

In order to ensure stationarity of the model economy, we further assume that commodity country productivity shock are cointegrated with the producer country shocks. Specifically, we assume that their cointegrating residual

$$q_t = \log z_{pt} - \beta \log z_{ct}$$

is stationary, following a mean-reverting jump-diffusion process

$$dq_t = [(1-\beta)(\mu - \mu_Z \eta) - \beta \psi q_t]\, dt + \sigma_p dB_{pt} - \beta \sigma_c dB_{ct} + dQ_t,$$

so that the commodity country productivity shock process (in logs) follows

$$d \log z_{ct} = (\mu + \psi q_t) dt + \sigma_c dB_{ct},$$

and therefore we can write

$$\begin{aligned} \frac{dz_{ct}}{z_{ct^-}} &= \left(\mu + \psi q_t + \frac{1}{2}\sigma_c^2\right) dt + \sigma_c dB_{ct} \\ &\mathrel{\hat{=}} \mu_{ct} dt + \sigma_c dB_{ct}. \end{aligned}$$

This cointegrated relationship can be interpreted as a reduced form representation of an economy where supply of the commodity is inelastic in the short run (based on the currently explored oil fields, say) but adjusts in the long run to meet the demand by the final good producers (e.g., as new fields are explored more aggressively when oil prices are high).

Similarly, we assume that shipping sector productivity is cointegrated with the commodity supply, with the cointegrating residual defined

$$q_{kt} = \log z_{ct} - \log z_{kt},$$

which follows a mean-reverting process

$$dq_{kt} = (\psi q_t - \psi_k q_{kt}) dt + \sigma_c dB_{ct} - \sigma_k dB_{kt}$$

so that the shipping shock process follows

$$\begin{aligned} d\log z_{kt} &= (\mu + \psi_k q_{kt})dt + \sigma_k dB_{kt} \\ \Rightarrow \frac{dz_{kt}}{z_{kt^-}} &= \left(\mu + \psi q_{kt} + \frac{1}{2}\sigma_k^2\right) dt + \sigma_k dB_{kt} \\ &\doteq \mu_{kt} dt + \sigma_k dB_{kt}, \end{aligned}$$

where the Brownian motions B_{pt}, B_{ct}, and B_{kt} are independent. The latter assumption captures the idea that shipping capacity cannot be adjusted quickly in response to shocks, which can lead to substantial volatility in costs of shipping over time, and therefore shipping costs that are very sensitive to demand shocks in the short run (e.g., Kalouptsidi (2011), Greenwood and Hanson (2013)). Our modeling of cointegrated jump-diffusion processes is similar to the model of cointegrated consumption and dividend dynamics in Longstaff and Piazzesi (2004). We can solve for output and commodity price dynamics by application of Ito's lemma (see Appendix).

2.3 Complete markets and consumption risk sharing

In order to emphasize that our mechanism does not rely on any financial market imperfections, we consider consumption allocations under complete markets. This is a standard benchmark in international finance, and is reasonable at least when applied to developed countries.[4] Under complete markets, the equilibrium allocation is identical to that chosen by a central planner for a suitable choice of a (relative) Pareto weight λ.

The planner's problem is therefore

$$V(z_{ct}, z_{pt}, z_{kt}) = \max_{\{X_t\}} E\left[\int_t^{\infty} e^{-\rho(s-t)} \left(\frac{c_{cs}^{1-\gamma} - 1}{1-\gamma} + \lambda \frac{c_{ps}^{1-\gamma} - 1}{1-\gamma}\right) ds \middle| \mathcal{F}_t\right],$$

where X_s is exports of final good to the commodity country, the commodity country consumption is $c_{cs} = X_s(1 - \tau_f(X_s, z_k))$, and the producer country consumption is $c_{ps} = y_{ps} - X_s$.

[4] For example, Fitzgerald (2012) estimates that risk-sharing via financial markets among developed countries is nearly optimal, while goods markets trade frictions are sizeable.

The first-order condition implies that

$$g(X_t, z_{ct}, z_{pt}, z_{kt}) \equiv \left[X_t(1 - \kappa_0^f - \kappa_1^f \frac{X_t}{z_{kt}})\right]^{-\gamma} \left(1 - \kappa_0^f - 2\kappa_1^f \frac{X_t}{z_{kt}}\right) - \lambda(y_{pt} - X_t)^{-\gamma} = 0 \quad (1)$$

must hold state by state for all t. In general, this nonlinear equation must be solved numerically, except for the special case of log utility ($\gamma = 1$).

Since the trade costs are increasing in the amount of goods shipped (holding shipping capacity fixed), the cost of transporting an extra unit of the final consumption good is increasing in total output y_{pt}. When output is high, the social planner allocates greater amounts of the good to the commodity country while shipping becomes increasingly costly.[5] The effects of individual state variables on the final good trade cost τ_f are displayed in Figure 1 as functions one shock while holding all other shocks constant at a value of 1.3. These effects are intuitive: greater shipping capacity decreases the cost of shipping, while higher productivity of the final goods producer increases trade costs by raising output and, consequently, the amount of goods shipped to the commodity country (higher productivity in the commodity country has a similar effect, as it feeds into final good output).

2.4 Exchange rates

The spot exchange rate in the absence of arbitrage is proportional to the ratio of the marginal utilities of the two representative agents,

$$\begin{aligned} S_t &= \lambda \frac{\pi_{pt}}{\pi_{ct}} = \lambda \left(\frac{c_{ct}}{c_{pt}}\right)^{\gamma} = \lambda \left(\frac{X_t(1 - \tau_f(X_t, z_{kt}))}{y_{pt} - X_t}\right)^{\gamma} && (2) \\ &= \lambda \left(\frac{1 - \tau_f(X_t, z_{kt})}{1 - x_t}\right)^{\gamma} = \left(1 - \kappa_0^f - 2\kappa_1^f \frac{X_t}{z_{kt}}\right) && (3) \end{aligned}$$

where the last equality follows from (1), implying that the real exchange rate is proportional to the marginal value to the commodity country consumer of a unit of the consumption good shipped from the country where it is produced (e.g., see Dumas (1992), Hollifield and Uppal

[5] The share of final good output that is exported can be increasing or decreasing in output, depending on the curvature of the utility function and the steepness of the trade cost profile: if the utility function is sufficiently concave, the planner compensates the increasing losses due to rising trade costs by increasing export share in good times (the empirically relevant case); otherwise, the share declines to reduce the deadweight loss.

Figure 1: Effect of Shocks on Shipping Costs

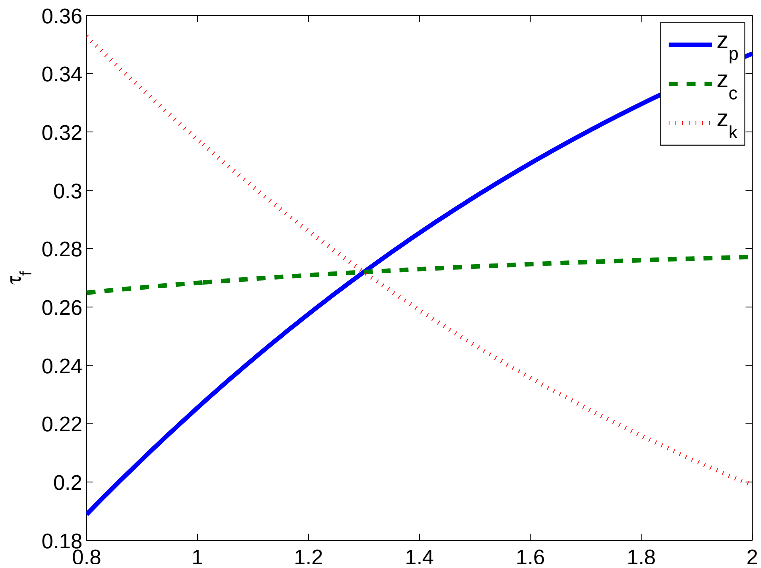

(1997), Verdelhan (2010)).

The real exchange rate is monotonic in the ratio of the two countries' consumption levels, is linear in the quantity of final good output exported to the commodity country, X_t, and is therefore closely related to the trade costs. Following good productivity shocks in either final good or commodity producing countries, total output y_p and exports X both increase, and therefore the producer country exchange rate depreciates. This is due to the fact that shipping costs lower the value of a marginal unit of the consumption good exported by its producer to the commodity country consumer, and more so when more of the good is shipped. Consequently, as (2) shows, both consumption and its marginal utility declines more slowly for the commodity country consumer than for the producer country consumer in good times, and also rises more slowly in bad times.[6] Positive shocks to the shipping capacity z_k reduce the cost of shipping and therefore act in the opposite direction, increasing the value of the unit of X to the commodity country and therefore lowering its exchange rate (X will

[6]In autarky, the commodity currency appreciates following good shocks to the final good production technology as its good becomes more highly demanded - this is the terms-of-trade effect, which is present even in the absence of complete financial markets, as emphasized by Cole and Obstfeld (1991). The effects of the commodity country productivity differ, however: terms of trade logic implies that commodity currency appreciates when the commodity becomes scarce following a bad supply shock. This is not generally true in our complete markets setup, as a decline in commodity supply leads to lower output of the final good, and higher value for the producer country currency.

Figure 2: Shocks and Exchange Rates

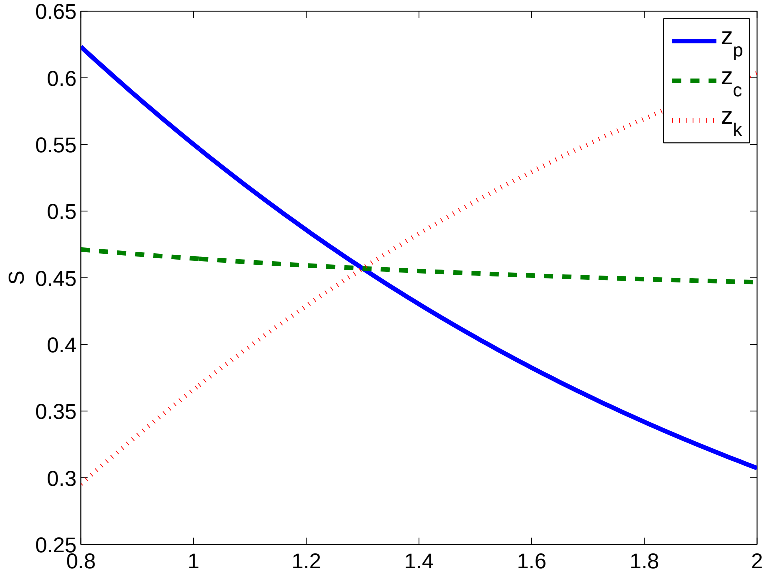

increase endogenously in response to higher shipping capacity, however, partially offsetting the influence of shipping cost shocks on the exchange rate.). These effects are displayed in Figure 2, which plots the exchange rate S (in units of commodity currency per one unit of final good producer currency), as a function of the three shocks, holding the other shock constant at a value of 1.3.

2.5 Asset pricing

Stochastic discount factors for the two countries are given by

$$\begin{aligned} \pi_{pt} &= e^{-\rho t} c_{pt}^{-\gamma} \\ \Rightarrow \frac{d\pi_{pt}}{\pi_{pt^-}} &= -\left\{\rho + \gamma\mu_{cpt} - \frac{1}{2}\gamma(1+\gamma)\sigma_{cpt}^T\sigma_{cpt}\right\} dt - \gamma\sigma_{cpt}^T dB_t + \left(e^{-\gamma \mathcal{J}_p} - 1\right) dN_t \end{aligned}$$

for the final good producer and

$$\pi_{ct} = e^{-\rho t} c_{ct}^{-\gamma}$$
$$\Rightarrow \frac{d\pi_{ct}}{\pi_{ct^-}} = -\left\{\rho + \gamma\mu_{cct} - \frac{1}{2}\gamma(1+\gamma)\sigma_{cct}^T\sigma_{cct}\right\} dt - \gamma\sigma_{cct}^T dB_t + \left(e^{-\gamma\mathcal{J}_c} - 1\right) dN_t$$

for the commodity producer, where $\mathcal{J}_p$ and $\mathcal{J}_c$ are log changes in the marginal utilities induced by jumps.

Risk-free rates are the (negative) drifts of the stochastic discount factors:

$$r_{pt}^f = \rho + \gamma\mu_{cpt} - \frac{1}{2}\gamma(1+\gamma)\sigma_{cpt}^T\sigma_{cpt} - \eta\mathbb{E}_Z\left[e^{-\gamma\mathcal{J}_p} - 1\right]$$

and

$$r_{ct}^f = \rho + \gamma\mu_{cct} - \frac{1}{2}\gamma(1+\gamma)\sigma_{cct}^T\sigma_{cct} - \eta\mathbb{E}_Z\left[e^{-\gamma\mathcal{J}_c} - 1\right],$$

for the final goods and commodity producer, respectively. The terms $\mathbb{E}_Z$ denote expectations taken over the distribution of jump sizes conditional on a jump occurring. The first two terms of the interest rate expressions above are equal between the two countries on average, as long-run consumption growth rates are equalized by the social planner. However, the last terms – the precautionary saving demands – differ. Since the final goods producer absorbs the bulk of productivity shocks to output, consuming a greater share in good times and a lower share in bad times, it experiences greater consumption volatility. Consequently, it has a greater precautionary demand and a lower interest rate on average. Similarly, the conditional expectation of marginal utility growth upon a jump is greater for the producer country consumer due to the same effect.

Since trade costs are persistent as long as shipping capacity adjusts slowly in response to demand, the interest rate variation is driven in part by the expected convergence in consumption due to cointegration (captured by the drift terms) and by the dispersion in conditional risk exposures of the pricing kernels (captured by the precautionary and jump terms). In particular, when output outstrips shipping capacity, the dispersion between the risk terms in the two countries is high, where as when shipping capacity is abundant relative to output this dispersion is lower. Figure 3 illustrates this effect for the case of logarithmic utility

Figure 3: Trade Costs and Endogenous Segmentation in Risk

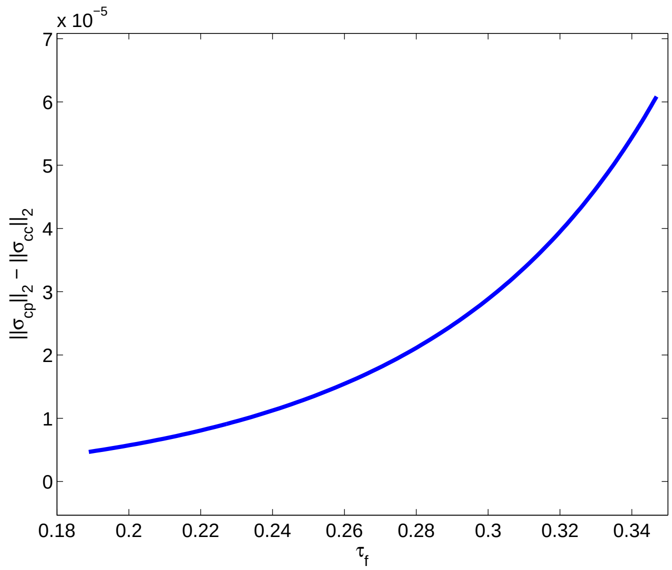

and no jumps: the difference between conditional consumption volatilities increases following good productivity shocks (or bad shipping sector shocks), which increase trade costs and consequently the degree of goods markets segmentation, reducing risk-sharing opportunities.

2.6 Expected excess returns: the carry trade

We can define the instantaneous excess return process for the currency trading strategy that is long the commodity currency (and short the producer currency) as

$$dRet_t = (r^f_{ct} - r^f_{pt})dt - \frac{dS_t}{S_{t^-}}.$$

This return can be earned by a final-good producing country investor directly, by shipping a unit of consumption good (borrowed at rate r^f_{pt}) and purchasing S_t units of the commodity-country risk free bonds, earning interest r^f_{ct} on these bonds, and converting it back into its own consumption good by shipping fewer units of the consumption good to the commodity country. It can also be obtained indirectly, by trading a state-contingent claim that replicates the payoff on this strategy, given complete financial markets. A commodity country investor can obtain a similar return, adjusted for the exchange rate.

The conditional expected excess return on this strategy (i.e., the currency risk premium) is given by the covariance of the exchange rate with the producer country pricing kernel (e.g., Bakshi and Chen (1997)):

$$E\left[dRet_t|\mathcal{F}_t\right] = E\left[\frac{dS_t}{S_{t^-}}\frac{d\pi_{pt}}{\pi_{pt^-}}|\mathcal{F}_t\right],$$

since the returns are expressed in the producer country numeraire (an equivalent statement holds for the commodity country pricing kernel if the returns are expressed in the commodity currency units). In general, this risk premium is not equal to zero, so that the uncovered interest parity relation $E\left[\frac{dS_t}{S_{t^-}}|\mathcal{F}_t\right] = (r^f_{ct} - r^f_{pt})dt$ need not hold.

In fact, this commodity currency trading strategy is profitable, on average, since the commodity currency is risky: it tends to appreciate in good times (when final good output is high) and depreciate in bad times, so that $E\left[dRet_t\right] = E\left[\frac{dS_t}{S_{t^-}}\frac{d\pi_{pt}}{\pi_{pt^-}}\right] > 0$. As long as exchange rates are persistent and close to random walks, the bulk of average carry excess return comes from the interest rate differentials. These effects are demonstrated in Figure 4, which plots sample paths of the key variables simulated from the model. While both interest rates fluctuate, with the commodity country interest rate being more volatile, and sometimes falling below that of the final good producer, on average the latter is lower. Therefore, a long position in the commodity currency and a short position in the "safe haven" currency of the final good producer is indeed a carry trade strategy, at least unconditionally. This strategy is a form of unconditional carry-trade strategy insofar as the commodity currency interest rate is on average higher than the producer country interest rate, i.e. as long as the precautionary terms are large enough.

Consistent with intuition, commodity currency exchange rate comoves with the commodity price P as well as realized shipping costs measured by $\tau_f(X, z_k)$ (for S the relationship is inverse). Interestingly, while carry trade returns are positively correlated with these variables, so are *expected* returns on the carry trade. This is due to the fact that the degree of dispersion between the conditional expected marginal utilities (and therefore the risk premium) is pro-cyclical, as trade costs are high in good times (especially if shipping capacity is lagging behind). The qualitative effect of trade costs on conditional volatilities of consump-

Figure 4: Model Dynamics: Example

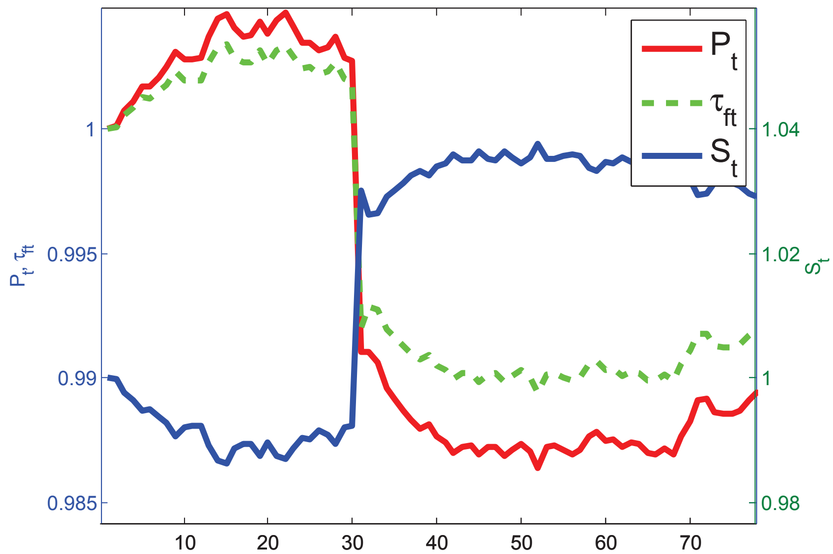

tion growth, which drive the risk premium in the absence of jumps, is displayed in Figure 3 above. We explore this mechanism quantitatively using the fully-specified model in Section 4 below.

2.7 Summary of implications

The model makes a set of predictions for the risk and return properties of exchange rates.

1. The final good-producing country bears more aggregate consumption risk. Therefore, it has a larger precautionary demand and lower interest rates, on average, than the commodity-producing country.

2. The commodity country currency is risky, as it appreciates in good times and depreciates in bad times. Therefore, it earns a risk premium, giving rise to a carry trade.

3. The commodity currency exchange rate (and therefore the carry trade) is positively correlated with the commodity price as well as the realized shipping costs, since they both increase in good times.

4. As high shipping costs imply lower degree of international risk sharing and therefore greater dispersion between conditional volatilities in consumption, conditional expected carry trade returns are positively correlated with trade costs.

Our model of exchange rate determination is deliberately simple and meant to highlight the mechanism leading to a carry trade: specialization combined with non-linear shipping costs. The model nevertheless makes a rich set of qualitative predictions, which we evaluate empirically before proceeding to analyzing its quantitative implications.

3 Empirical evidence

3.1 Data

Following Lustig, Roussanov, and Verdelhan (2011) we use forward and spot exchange rates to construct forward discounts (approximately equal to the interest rate differentials by the covered interest parity relation) and excess returns on currencies. We use the same set of currencies. Data is provided by Barclays and Reuters and is available via Datastream. We use monthly series from January 1988 to December 2012.[7]

We use two samples in our analysis. The sample of all 35 developed and emerging countries includes: Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Euro area, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, Italy, Japan, Kuwait, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, United Kingdom. The sub-sample of 21 developed-country currencies includes: Australia, Austria,

[7] While Lustig, Roussanov, and Verdelhan (2011) start their sample in 1983, very few currencies have forward discounts available in the first few years of the sample, as a number of countries, including Australia and New Zealand, undergo transition from fixed to floating exchange rates during this period. The latter countries have forward discounts available starting in 1985, but these display patterns suggesting episodes of extreme illiquidity, such as large bid-ask spreads and violations of covered interest parity relation (CIP) before 1988. Finally, the Plaza Accord of September 22, 1985 led to a large but gradual appreciation of the Deutschmark, the French Franc, and the Japanese Yen over the course of 1986 and 1987. Since these movements were largely predictable by investors it appears natural to consider unconditional strategies including these currencies starting in 1988.

Belgium, Canada, Denmark, Euro, Finland, France, Germany, Greece, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom.

Table 1 shows U.S. dollar average returns and forward discounts on the nine most actively traded currencies, collectively known as the G10 countries (the tenth currency being the U.S. dollar itself), over our sample period. The German Deutschmark forward discount and the excess return to investing in Deutschmark forward contracts prior to 1999 are spliced with the euro variables post-1999. The table is sorted from low average returns to high average returns. What is immediately apparent is that the high return countries tended to have unconditionally high forward discounts, consistent with the unconditional carry trade strategy documented in Lustig, Roussanov, and Verdelhan (2011).

Table 1: **G10 Currency Average FX Returns and Discounts**

Country	Excess Return	Forward Discount
Japan	-1.97	-2.70
Switzerland	-0.32	-1.53
Germany/Euro	0.11	-0.15
Sweden	0.80	1.37
United Kingdom	0.92	1.81
Canada	1.66	0.65
Norway	1.99	1.81
Australia	4.02	2.71
New Zealand	4.06	3.08

Average annualized forward discounts and excess returns (without accounting for transaction costs) for the "G-10" currencies from the perspective of a U.S. dollar investor. Germany/Euro is calculated based on the German Deutschmark prior to 1999 and the Euro post 1999. Data are monthly futures from 1988 to 2012 taken from Datastream.

Interestingly, this relation between average forward discounts and excess returns is not a perfectly monotonic one, in that some low return countries have high discounts. This is not necessarily surprising since factors other than expected returns (e.g. expected inflation) can have an effect on nominal interest rates, and therefore forward discounts.[8] It is clear, however, that the countries with low returns tend to be countries with advanced manufacturing

[8]Pairwise average currency returns are only marginally statistically different from zero due to the substantial noise in bilateral exchange rate movements, consistent with evidence in Bakshi and Panayotov (2012); however, aggregating currencies into portfolios (e.g., long bottom four, short top four) reduces idiosyncratic noise and ensures robustly statistically significant average returns (as detailed in Data Appendix Table A-1).

economies which are also relatively resource poor. Indeed, the entire top half of the table: Germany, Japan, Sweden, Switzerland, and the UK all fit this description to some degree. In contrast, the high return countries on the bottom half of the table tend to be large exporters of either oil (Canada and Norway) or other base agricultural or mineral commodities (New Zealand and Australia). This simple observation suggests a potential unconditional carry trade strategy based on the trade characteristics of each country.

In order to classify countries based on their exports we utilize the U.N. COMTRADE database of international trade flows. We use the NBER extract version of this data, available for years 1980-2000, we augment it with the original COMTRADE data for years 2001-2012 following the same methodology. The two goods in the model are a basic good, which is used as an input in production, and a final good, which is used in consumption. While this suggests a potential classification of goods as either "input" or "final" goods, there are many goods for which this classification struggles to conform to the intuition of the model. The important mechanism in the model hinges on the extra trade costs associated with shipping complex produced goods back to the commodity exporter rather than the specific use of the goods as consumption or input. For instance, New Zealand is a large exporter of many agricultural commodities, some of which (such as butter) are in their final consumable form. Likewise, New Zealand imports a large amount of sophisticated construction equipment which is produced using basic commodities (e.g., metals, energy) as an input. However, in the context of the model, a complex piece of construction equipment seems more closely related to the final good rather than the basic good, while butter is a better representation of the basic good. Therefore to be consistent with the model mechanism we classify goods as a basic good (i.e. a commodity) or a complex good based on their 4-digit SITC codes. The classifications at the 2-digit level are in the appendix (Table A-2), and the full classification is available upon request.

Using this classification of goods we create two different country-specific measures, the first is the ratio of each countries' net *exports* in basic goods to its total trade in basic goods in each year of the formation period, and the second is the ratio of net *imports* in complex goods to its total trade in complex goods. Both of these measures by construction take a value between -1 and 1. The first sort captures the extent to which a country specializes in the production of basic commodities, and the second variable captures the extent to which a country imports complex goods. Intuitively, for a given country a high ratio of commodity exports tends to be accompanied by a high ratio of complex imports.

3.2 Currency portfolios sorted on Import/Export data

The main prediction of the model is that countries exporting basic goods and importing complex goods should have lower exposure to global productivity shocks, and therefore their currencies should have higher average discounts and earn higher returns. Figure 5 plots the average forward discounts on individual currencies over the time period following the creation of the euro (post-1999) against the average ratio of the final good *imports* plus basic good *exports* to total trade over the whole sample (1988 to 2012). The two variables appear to line up well, with higher levels of the import ratio typically corresponding to high average forward discounts (e.g., this includes the so-called "commodity countries" - Australia, New Zealand, Norway, and South Africa), where as low values of final good import ratio correspond to low average forward discounts (Japan is the most salient extreme case). The exceptions to the pattern tend be countries experiencing high inflation over the sample period (Mexico, Hungary, and Philippines).

Figure 5 also displays a cross-sectional regression line that relates our import/export composition variable to the mean forward discounts. As indicated by the R^2 of this regression, our trade-based variable explains a third of the cross-sectional variation in the average interest rate differentials across countries. This variation is clearly not driven entirely by country size as suggested by Hassan (2010), since the U.S. as well as the U.K. are in the middle of the distribution of the import/export variable (as well as of the average forward discount, which equals zero for the U.S. by construction)[9] These distributions are consistent with evidence in Lustig, Roussanov, and Verdelhan (2013) who estimate that the U.S. as well as the U.K. pricing kernels have approximately average loadings on the global factor that gives rise to

[9]Hassan (2010) shows a relation between country size and and currency risk premia using panel regressions of returns and forward discounts on GDP. We perform similar regressions and find that both trade ratios and GDP have significant explanatory power for expected returns and forward discounts, with the effects being slightly stronger for import ratios, particularly in the latter half of the sample (See Table A-4 in the Data Appendix). However, neither variable completely drives out the other, so we view the mechanisms as complementary.

the unconditional currency risk premia (among developed countries), where as Japan and Australia are at the opposite ends of the spectrum.

In order to examine the patterns of average excess returns predicted by the model, we sort all of the countries in our sample into 6 portfolios (5 for the subsample of developed countries) using the rolling five-year average of the export ratio of input goods. We then repeat this strategy using the import ratio of complex goods. We compute the average forward discounts and average log excess returns for each of the portfolios.

Average forward discounts and average returns are computed from 1988-2012.[10] The construction of these portfolios represents an implementable trading strategy, relying only on trade data from available at the time of portfolio formation. Furthermore, since the composition of countries' imports and exports is generally stable over time, the strategy is essentially an unconditional carry-trade strategy, similar to the unconditional interest rate strategy described by Lustig, Roussanov, and Verdelhan (2011).

We work with one-month forward and spot exchange rates in units of foreign currency per U.S. dollar, denoted by F_t and S_t, respectively. Using the individual currency one-month forward discounts $f_t - s_t$ (lower case letters representing logarithms) and log excess returns approximated as

$$rx_{t+1} = f_t - s_{t+1},$$

we compute the log currency excess return rx^j_{t+1} for each portfolio $j = 1, 2 \ldots, 6$ by averaging over N_j currencies in the portfolio:

$$rx^j_{t+1} = \frac{1}{N_j} \sum_{i \in N_j} rx^i_{t+1}. \tag{4}$$

Similarly, currency portfolio excess returns (in levels) RX^j are computed by averaging individual currency excess returns in levels, $RX^i = (F^i_t - S^i_{t+1})/S^i_t$ analogously to (4). We do not take into account bid/ask spreads in the construction of these portfolios at the monthly frequency. Since our portfolios require very little rebalancing, transaction costs are likely to

[10]Currency forward data is available starting from 1983, but only for a few currencies, and from 1985 for most of the developed country currencies in our sample. In order to construct the portion of the standard currency carry-trade unrelated to the commodity-currency carry-trade constructed using import and exports we rely on three year rolling regressions, resulting in a post 1988 sample period. Details are in Section 3.3.

Figure 5: Average Forward Discounts and Import Ratios

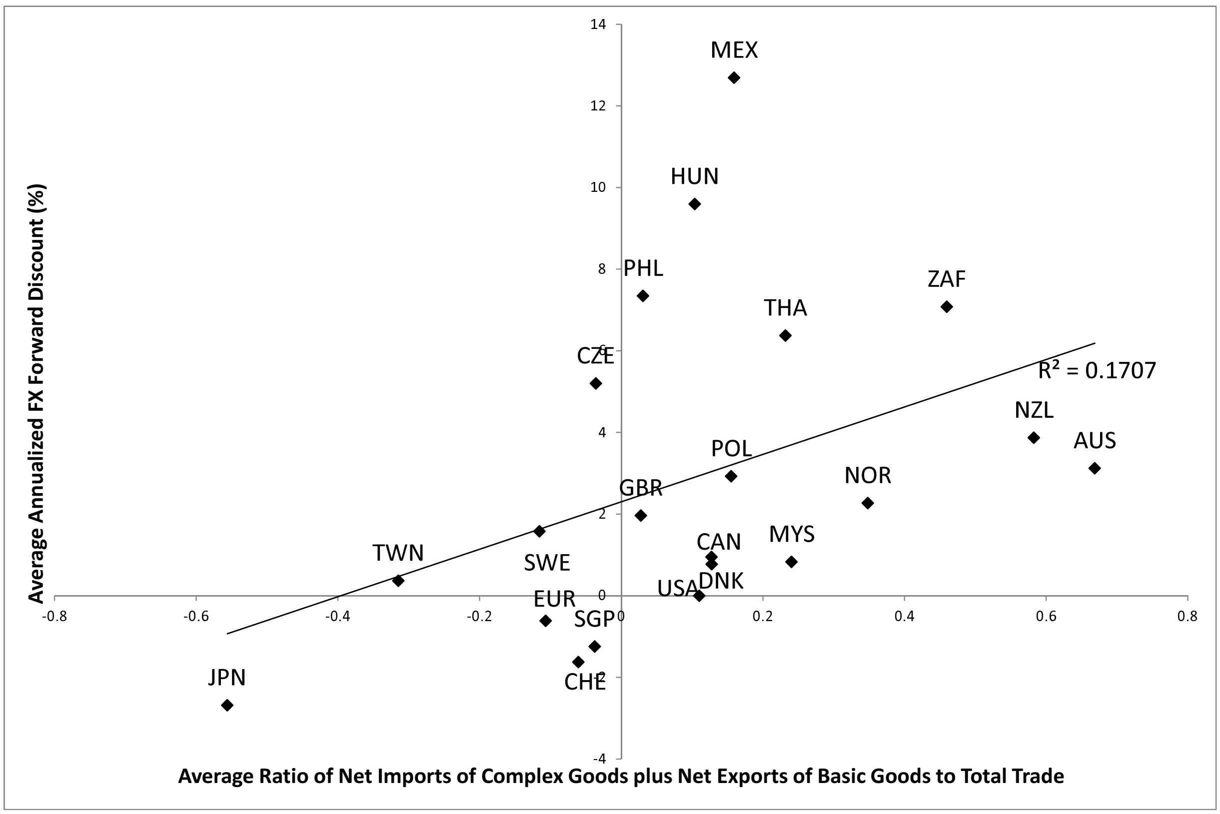

This figure plots average forward discounts from 1988 to 2012 against a combined measure of the extent to which a country exports basic goods and imports complex goods. The measure is constructed by adding the net *imports* of complex goods plus net *exports* of basic goods and then dividing by total trade in all goods. This ratio is calculated in each year and averaged over the 1988 to 2012 sample for each each country. The FX discount for the German Deutschmark and the Euro are treated a single observation and are plotted against the import ratio for the Eurozone. Trade data are annual, from UN Comtrade (available via NBER extracts from 1980 through 2000), while spot and forward exchange rate data are monthly, from Barclays and Reuters (available via Datastream).

Table 2: Currency Portfolios Sorted on Final Good Exports

Portfolio	1	2	3	4	5	6	1	2	3	4	5
	Panel I: All Countries						Panel 2: Developed Countries				
	Forward Discount: $f^j - s^j$						$f^j - s^j$				
Mean	-0.79	1.43	2.15	1.66	2.52	2.76	-1.04	0.23	0.54	1.32	2.89
Std	0.60	0.70	0.95	0.53	0.60	0.52	0.64	0.77	0.73	0.61	0.50
	Excess Return: RX^j						RX^j				
Mean	0.10	1.90	3.73	2.24	2.29	4.67	0.56	1.74	2.47	1.84	5.55
Std	7.90	8.65	10.19	7.57	8.38	9.68	8.99	10.58	9.07	8.66	10.62
SR	0.01	0.22	0.37	0.30	0.27	0.48	0.06	0.16	0.27	0.21	0.52

This table reports average forward discounts and average log excess returns on currency portfolios sorted on the ratio of the countries' net exports of finished goods relative to total trade in such goods, in descending order. Each year's ranking is computed using the average ratio for the prior four years. Trade data are annual, from UN Comtrade (available via NBER extracts). Forward and spot exchange rate data are monthly, from Barclays and Reuters (available via Datastream). The returns do not take into account bid-ask spreads. The sample period is 1/1988 to 12/2012.

be small (returns based on long-horizon, e.g. one-year, forward contracts are typically similar to those obtained by rolling over shorter-horizon contracts; we report the results using one-year forward contracts with bid-ask spreads in the Data Appendix.).[11]

The results are reported in Tables 2 and 3. The results using both sorts are very similar: portfolios representing high complex good export ratios and those with high basic good export ratios have low average forward discounts, suggesting that they capture countries whose interest rates are typically low relative to the U.S. Conversely, portfolios with high values of the commodity exports ratio and low values of final good exports exhibit high average forward discount, indicating high average interest rates. The pattern is virtually monotonic across portfolios for both sorts, especially for developed countries subsample, with differences between the highest and the lowest portfolios' average forward discounts of around 4% per annum for the basic good sort over 5% per annum for the complex good sort.

Importantly, portfolio average excess returns follow the pattern of the average forward discounts, being negative for the low portfolios and positive for the high portfolios, with the

[11]The portfolio is rebalanced to handle the introduction of the Euro. Prior to 1999 breakpoints are calculated including the component countries of the Euro as separate entities. Post 1999 the breakpoints are recalculated counting the Eurozone as a single country.

Table 3: Currency Portfolios Sorted on Basic Good Exports

Portfolio	1	2	3	4	5	6	1	2	3	4	5
	Panel I: All Countries						Panel 2: Developed Countries				
	Forward Discount: $f^j - s^j$						$f^j - s^j$				
Mean	-0.37	0.53	0.94	3.89	2.78	2.01	-0.95	0.29	0.90	1.63	2.16
Std	0.66	0.62	0.61	0.86	0.56	0.48	0.73	0.58	0.74	0.61	0.58
	Excess Return: rx^j						rx^j				
Mean	0.12	2.22	1.69	4.43	2.68	3.76	0.65	0.93	2.21	4.22	4.04
Std	7.89	9.16	8.71	9.14	7.83	8.81	9.01	9.57	10.51	8.97	9.33
SR	0.02	0.24	0.19	0.48	0.34	0.43	0.07	0.10	0.21	0.47	0.43

This table reports average forward discounts and average log excess returns on currency portfolios sorted on the ratio of the countries' net exports of basic input goods relative to total trade in such goods, in ascending order. Each year's ranking is computed using the average ratio for the prior four years. Trade data are annual, from UN Comtrade (available via NBER extracts). Forward and spot exchange rate data are monthly, from Barclays and Reuters (available via Datastream). The returns do not take into account bid-ask spreads. The sample period is 1/1988 to 12/2012.

spreads in average returns between extreme portfolios close to 4% per year for both the basic good sort and the complex good sort. Thus, the differences in the average forward discounts translate almost fully into average excess returns, contrary to the UIP hypothesis. Since the sorting variables are very persistent, these differences are likely to capture unconditional rather than conditional risk premia.

3.3 Comparison with traditional carry trade strategies

To facilitate comparison with traditional carry-trade strategies, we sort countries based on a measure of the extent to which the country both exports basic goods and imports complex goods, constructed as the sum of net exports of basic goods and net imports of complex goods, divided by the total trade in all goods. Average forward discounts and excess returns for these portfolios are shown in Table 4. We then consider returns on a portfolio which is long the portfolio with the highest ratio and short the lowest among all countries over the prior four years. We refer to this strategy as IMX (*Importers minus eXporters* of finished goods). We then construct two additional carry-trade strategies. The first uses the traditional method

Table 4: Currency Portfolios Sorted on Combined Imports/Exports Measure

Portfolio	1	2	3	4	5	6	1	2	3	4	5
	Panel I: All Countries						Panel 2: Developed Countries				
	Forward Discount: $f^j - s^j$						$f^j - s^j$				
Mean	-0.39	0.84	1.70	2.01	2.91	2.78	-0.68	-0.08	0.91	1.48	2.44
Std	0.64	0.68	0.63	0.69	0.61	0.49	0.68	0.73	0.58	0.65	0.54
	Excess Return: rx^j						rx^j				
Mean	-0.01	2.15	3.38	2.19	2.44	4.96	-0.12	1.01	0.99	3.48	3.86
Std	7.93	9.06	9.66	7.08	8.75	9.65	8.98	10.66	9.26	8.89	9.63
SR	0.00	0.24	0.35	0.31	0.28	0.51	-0.01	0.09	0.11	0.39	0.40

This table reports average forward discounts and average log excess returns on currency portfolios sorted on a ratio designed to capture the extent to which each country exports basic goods and imports finished goods. The ratio is constructed by adding the level of net exports in basic goods to the level of net imports in finished goods, and then dividing by the level of total trade in all goods. Each year's ranking is computed using the average ratio for the prior four years. Trade data are annual, from UN Comtrade (available via NBER extracts). Forward and spot exchange rate data are monthly, from Barclays and Reuters (available via Datastream). The returns do not take into account bid-ask spreads. The sample period is 1/1988 to 12/2012.

of sorting currencies based on the interest rate. Following Lustig, Roussanov, and Verdelhan (2011) we follow a strategy forming portfolios based on the current interest rate in each month, and label this strategy HML_{FX}. In addition, in order to construct a strategy which is related to the part of the standard carry trade not related to the IMX strategy, we construct a tradeable strategy that is long HML_{FX} and short a number of units of IMX equal to its contribution to HML_{FX}. This strategy (which we refer to as $CHML_{FX}$) is calculated as $CHML_{FX,t+1} = HML_{FX,t+1} - \beta_t^{HML,IMX} IMX_{t+1}$, where $\beta_t^{HML,IMX}$ is estimated using a 3-year rolling regression up to time t.

Table 5 reports the returns and standard deviations of the portfolios for each of these strategies. By construction $CHML_{FX}$ and IMX have very low correlation, while both strategies are positively correlated with HML_{FX}. While the import-based strategy underperforms the traditional carry trade strategy, it does have a significantly positively return. Brunnermeier, Nagel, and Pedersen (2008) suggest that crash-risk is important for understanding carry-trade risk, interestingly this table shows that the portion of the traditional carry-trade

Table 5: Carry-Trade Strategies

Strategy	Mean	St. Dev.	SR	Skewness	Correlation Matrix	
IMX	4.97	9.18	0.54	-0.53		
	(1.82)	(0.53)	(0.22)	(0.28)		
HML_{FX}	9.63	9.44	1.02	-0.36	0.26	
	(1.87)	(0.48)	(0.21)	(0.19)	(0.07)	
$CHML_{FX}$	8.20	9.01	0.91	0.13	-0.06	0.86
	(1.77)	(0.43)	(0.16)	(0.21)	(0.07)	(0.03)

Summary characteristics of returns on different carry-trade strategies. IMX is the return on a strategy long the currencies of complex good importers and short exporters, based on the combined imports/exports measure. Imports and exports are the average over a rolling window of the three prior years. HML_{FX} is the return on a strategy which is long high-interest rate countries and short low interest rate countries which is rebalanced each month. $CHML_{FX}$ is the return of a strategy which is long HML_{FX} and short a proportional amount of IMX where the proportion is determined using a 3-year rolling regression of HML_{FX} on IMX. The returns do not take into account bid-ask spreads. Bootstrap standard errors are in the parentheses.

related to IMX seems to account for nearly all of the negative skewness in traditional carry trade strategy.

3.4 Explaining the carry trade with IMX factor

While the high return of $CHML_{FX}$ shows that the IMX factor does not completely subsume the traditional carry trade, there appears to be a portion of carry-trade returns that is related to the characteristics of countries' trade, which are very stable over time. Again the magnitude of the return differential is similar to the unconditional interest rate carry-trade in Lustig, Roussanov, and Verdelhan (2011), who show that roughly half of the carry-trade premium can be explained as an unconditional premium on countries with a high average interest rate compared to those with a low average interest rate. To test if the import/export sort is capturing the same effects, we follow Lustig, Roussanov, and Verdelhan (2011) and construct an unconditional sort based on the average interest rates of countries over a preformation period from 1984 - 1995, and then examine portfolio returns over a period from 1995 to 2012. We term the return to this strategy $UHML_{FX}$. We then test whether the IMX factor can explain the positive returns to the traditional interest rate carry trade strategies, $UHML_{FX}$ and HML_{FX}.

Table 6 reports regressions of the form

$$RX_t^j = \alpha_j + \beta_j IMX_t + \epsilon_t^j,$$

where test assets i in the regression are the component portfolios (rebalanced according to interest rate each period) of both the standard HML_{FX} factor as well as the component portfolios (sorted based on the average interest rate over the pre-1995 formation period) of the $UHML_{FX}$ strategy, in addition to the long-short strategies HML_{FX} and $UHML_{FX}$. The results show that the IMX strategy fully explains the returns to the $UHML_{FX}$ strategy, with monotonically increasing betas, insignificant alphas, and high R^2, while explaining only some of the returns to the traditional HML_{FX}. These results emphasize that the mechanism in this model is most useful in understanding the returns to the unconditional portion of the carry trade, due the fact that the composition of traded goods for each country is highly stable through time. This is consistent with the evidence in Lustig, Roussanov, and Verdelhan (2013) who show that two separate global factors are needed to explain the unconditional and the conditional currency risk premia. It is not surprising that the IMX can explain the unconditional component of the currency carry trade, but not the conditional component, since there is much greater persistence in the countries import-export patterns than in their risk-free rates. The traditional carry trade captured by the HML_{FX} factor captures both the conditional and the unconditional risk premia, where as IMX only captures the latter, as predicted by the theory.[12]

To further shed light on the underlying mechanism, we now turn to the relation between carry-trade strategies and the salient variables of the model.

[12]Lustig, Roussanov, and Verdelhan (2011) show that accounting for transaction costs (in the form of bid ask spreads) can reduce the profitability of the traditional carry strategies. While our excess return definition does not account for transaction costs, the latter are unlikely to have a major impact on the profitability of our IMX strategy or the unconditional carry trade strategy, since it requires much less frequent rebalancing. We verify this by constructing annually rebalanced strategies with excess returns based on 12-month forward contracts using bid and ask quotes published by Reuters, which imply a Sharpe ratio for the IMX strategy of 0.42, as reported in the Data Appendix Table A-3.

3.5 Differences in risk exposure across countries

The model's key prediction is that commodity country consumption is less risky than that of the final good-producing country. While our two-country model is too stylized to be taken to the data directly, we provide evidence by grouping countries that more closely resemble the two types. We form two baskets of G10 currency countries, the four "commodity countries" of Australia, Canada, New Zealand, and Norway, and the four "producer countries" of Japan, Eurozone / Germany, Sweden and Switzlerland. Table 7 displays the standard deviation of quarterly consumption growth rates for the two baskets over the period 1993-2012. As the model predicts, aggregate consumption growth of final goods producers is more volatile than that of commodity producers (1.25% per annum vs. 0.88%).

The model predicts that producer country consumption is more sensitive to the global productivity shocks that are transmitted into the carry trade, rising faster in good times (when carry strategy does well) and declining in bad times (when carry trade does poorly). We can evaluate this prediction by computing the consumption betas for the commodity-currency carry trade factor IMX using both baskets. As indicated in Table 7, producer country consumption is almost twice as sensitive to the carry returns, compared to the commodity-country consumption, with IMX betas of 0.033 for the producer basket and 0.013 for commodity countries. The short sample makes for imprecise estimates, and the high volatility of the IMX factor relative to changes in consumption growth makes for low absolute magnitudes of the betas, but the final goods producers' consumption beta is significant at the 10% level using OLS standard errors (though not significantly different from the commodity countries betas).

Another important test of the mechanism in the model is the exposure of different countries' marginal utility to shocks to global productivity. The model predicts that wealth of commodity exporting (and complex good importing) countries should have lower exposure to global economic shocks and hence IMX. While we do not observe aggregate wealth holdings, we can attempt to approximate them using equity market wealth. To this end we collect country specific MSCI equity indices for 19 developed countries. For each country we perform

Figure 6: Stock Market Betas and Import Ratios

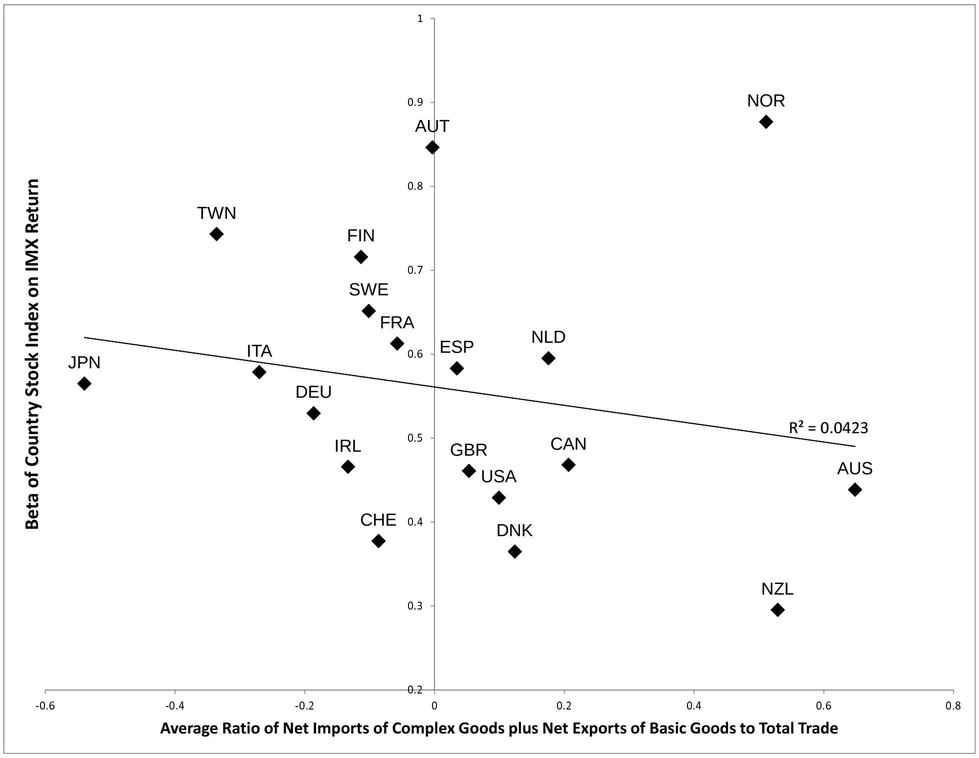

This figure plots the betas from the regression

$$R^e_{j,t} = \alpha^e_j + \beta^e_j IMX_t + e_{j,t}$$

where $R_{j,t}$ is the return to the market equity index for each developed country j. Betas are plotted against each country's combined import/export ratio measure as described in Table 4. Equity returns are from global financial data. Data is monthly from 1988 to 2012, from Datastream.

a regression of the return to the equity index on the return to IMX.

$$R^e_{j,t} = \alpha^e_j + \beta^e_j IMX_t + e_{j,t} \tag{5}$$

Figure 6 shows a graph of the β_j for each country as a function of the import-ratio of complex goods. The graph shows, that withe the notable exception of Norway, β_j tends to be a decreasing function of the import ratio. In other words, stock returns in countries which tend to be importers of these goods have less exposure to the innovations to IMX, again consistent with the predictions of the model.

3.6 Currency carry-trades, commodity prices, and shipping costs

In this section we examine the contemporaneous relation between the different carry-trade strategies and two of the important variables in our model: commodity prices and shipping costs. According to the model, if the returns to IMX are compensation for exposure to global economic activity, we should expect returns to this strategy to be positively correlated with changes in commodity prices. Since convex shipping costs in the model are the key drivers of the carry trade excess returns, we also expect that positive shocks to global productivity should increase trade costs while also generating a positive return to IMX. Therefore, we expect realizations of IMX to be positively correlated with changes in trade costs. To test the first hypothesis we use the Commodity Research Bureau's all commodity spot index. In order to proxy for levels of trade cost we use the Baltic Dry Index (BDI).

Table 8 reports contemporaneous regressions of the three currency carry-trade strategies on innovations to the logs of the CRB commodity index and the BDI over the whole sample. The IMX strategy loads heavily on these two variables, with contemporaneous R^2 near 15%. The traditional carry-trade loads on them as well, but the relationship is much weaker, and the residual component $CHML_{FX}$ has very little relation with these two variables with negligible R^2.

This is again consistent with the mechanism in the model. Since the composition of exports for a given country is very stable through time, we would expect the predictions of the model to explain an unconditional carry trade strategy but to be less likely to explain a strategy relying upon a continuous rebalancing of portfolios. We interpret the fact that these predictions of the model are only present in the unconditional portion of the carry trade strategy as evidence for this explanation.

Though the model does not distinguish between different types of commodities, it is interesting to see which commodities have the strongest correlation with the constructed IMX factor. Table 9 reports betas of different commodity sector indices from the CRB, as well as an index of energy commodities and several metals not included in the CRB indices. Commodities which are inputs into production, namely energy commodities, raw industrials, and industrial metals, tend to have the highest loadings on the IMX factor, again broadly consistent with the model's intuition.

3.7 Case study: the global financial crisis

As a further illustration of the model mechanisms in the data, we examine the behavior of model variables during the global financial crisis, which coincided with a dramatic decline in output, especially among final good producer countries, such as Japan, and a collapse in international trade volume (e.g. see Eaton, Kortum, Neiman, and Romalis (2011)). As Figure 7 shows, the data lines up nicely with the model predictions over this period. Panel A shows that the commodity currencies tended to depreciate relative to final good producer currencies during the crisis. Panel B illustrates that this is reflected in a large negative return on the IMX strategy, and that this return is accompanied by large negative changes in the CRB Commodity Spot Index and the Baltic Dry Index. Perhaps most importantly, even though commodity prices were dropping during this period, Panel C shows household consumption growth of the commodity countries did not fall as severely as that of the producer countries. The outliers are two of the smaller countries, New Zealand and Switzerland. Panel D shows that a GDP-weighted basket of commodity countries' consumption growth greatly outperforms that of final goods producers during the crisis.

3.8 Predicting carry-trade returns

In addition to contemporaneous correlations between carry returns, commodity prices, and the Baltic Dry Index, another important implication of the convex adjustment costs in the model is that the difference in exposure to the aggregate shock of the two countries is more severe during times when it is costly to ship goods, this leads a predictive relation between the level of shipping costs and expected return on the carry trade. Since shipping costs tend to be high in good times, when output and commodity prices are also high, this expected return is pro-cyclical.

This mechanism is consistent with recent evidence of carry trade return predictability with pro-cyclical variables, such as commodity prices, documented by Bakshi and Panayotov (2012). Similarly, Bakshi, Panayotov, and Skoulakis (2010) show that high levels of the BDI predict high returns in many different asset classes, including commodities. We document a similar predictive relation between the BDI and the traditional currency carry trade, but

Figure 7: Currencies, Commodities, Trade Costs, and Consumption During the Crisis

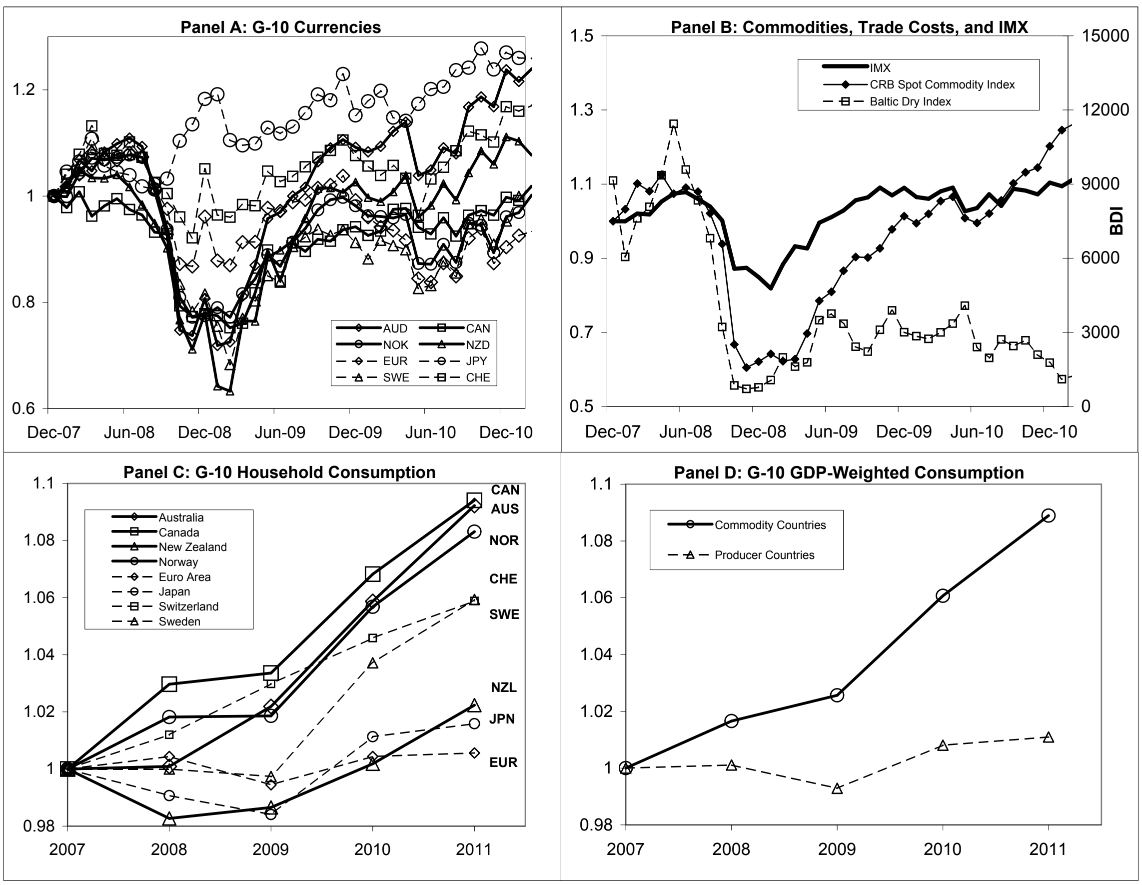

Currency and economic variables during the global financial crisis. Panel A shows monthly cumulative currency returns on the four G10 "commodity countries" (Australia, Canada, New Zealand, and Norway) and the four G10 "producer countries" (Europe, Japan, Switzerland, and Sweden). Panel B shows the monthly performance of the IMX strategy as well as monthly changes in the Commodity Research Bureau All Commodity spot index and the Baltic Dry Index (BDI). Panel C shows household consumption of the eight countries, and Panel D shows the consumption growth of GDP-weighted baskets of the two country groups. All exchange rate, commodity price, and consumption variables normalized to one in December 2007. Data from Datastream and the OECD.

find that all of the predictability is concentrated in the unconditional portion of the strategy, as captured by IMX. We also confirm the predictability of the traditional carry-trade in the G10 currencies Bakshi and Panayotov (2012) using lagged innovations in commodity spot indices, and again find that the predictability is concentrated in the portion of the trade related to our trade sorts.

To test for predictability we perform univariate predictive regressions analogous to Bakshi, Panayotov, and Skoulakis (2010)

$$rx_{i,t} = a_i + b_i \Delta bdi_{t-4,t-1} + \tilde{\epsilon}_{i,t} \tag{6}$$

Where i represents the four carry trade strategies, rx is log excess return over horizons $i \in \{1, 3, 6, 12\}$ months, and the predictive variable, $\Delta bdi_{t-4,t-1}$, is the change in the log of the BDI over the prior three months. Table 10 shows the results. We find a strong predictive relation in IMX, with R^2 of 4%. This relation is still significant but with lower R^2 for the standard carry-trade HML_{FX} factor. Most interestingly, the relation completely disappears when considering $CHML_{FX}$, which captures the portion of HML_{FX} that is orthogonal to IMX. Following Bakshi and Panayotov (2012), we repeat this exercise using innovations to the CRB Industrial Metals index and carry-trade strategies constructed in the G10 currencies, and see the same result (Table 11). To the extent there is predictability in the HML_{FX} carry trade return, it is primarily due to the predictability of the IMX portfolio. Again, the predictions of the model match nicely with the observed behavior of the unconditional strategy.

4 Quantitative analysis

So far we have only explored the qualitative implications of our model. We now turn to quantitative analysis. Ideally, we would like to calibrate the model parameters to closely match empirical moments. The fact that the model features only two countries (each completely specialized in producing one kind of good) makes such a moment-matching exercise challenging. In order to circumvent this challenge we make an assumptions that countries that are ranked at the top of the final good exporter measure as a group are representative of a final-good producer country in the model, while countries that rank at the bottom (i.e., the final good importers) are representative of the commodity country. Our empirical results above appear to corroborate this distinction, even though the difference between the two types of countries is much less stark in reality than our model assumes. We form two baskets using the set of G10 countries: one of the countries with the four highest import ratios (commodity countries) and the other of the four lowest (final good producer countries). We average macroeconomic and financial variables across countries within each basket and compare their properties to those implied by the model. Table 12 summarizes these moments while Table 13 lists the parameter values used in the calibration.

We present the summary statistics from the model-generated simulated data in three ways: we simulate the model 10,000 times, each time generating sample periods of approximately the same length as those in our data (30 years). Besides reporting both mean and median statistics across the simulations we also report means conditional on no "disasters" occurring in the sample (i.e. jumps that imply an annual consumption drop in the final good producer country that is greater than 5%). This definition is conservative, as Barro (2006) defines disasters as consumption drops of 10% and greater. We calibrate the distribution of jump sizes so that its tail approximately corresponds to the distribution of empirically observed consumption disasters compiled by Barro and Ursua (2008) (the largest disaster in their sample corresponds to a consumption drop of 70%, which is approximately the same as the upper bound of our jump distribution $Z_{max} = 1.2$). Disasters - large jumps that cause a 5% or greater drop in consumption - occur at least once over a 30-year period with probability of 16% in the simulated samples given that the jump intensity η is such that a jump occurs on average every 25 years, the smallest jump size is 2%, and the power law distribution of jump sizes has a tail exponent of 1.1.[13] Since the probability of such jumps is sufficiently small, these conditional statistics capture the sense in which rare disasters contribute to the observed risk premia. There is some debate in the literature about the extent to which rare disasters and peso problems contribute to currency risk premia[14]. While the economic

[13]Backus, Chernov, and Martin (2011) argue that equity option prices imply lower probabilities of consumption disasters than the magnitude required to match the equity premium.

[14]Models such as Farhi and Gabaix (2008) and Gourio, Siemer, and Verdelhan (2013) rely on rare disasters for explanations of the forward premium puzzle. Empirical evidence in Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2009), Jurek (2009), Burnside, Eichenbaum, Kleshchelski, and Rebelo (2008), and Chernov,

mechanism of our model does not rely on rare disasters, the simulation results reveal that the possibility of such disasters that may occur but are not observed in sample substantially improves the model's ability to quantitatively account for the carry trade risk premium that is generated by the spread between the higher-order moments of the marginal utilities in the two countries.

The modest degree of relative risk aversion $\gamma = 5$ ensures that the model does not overshoot the exchange rate volatility observed in the data too much in the absence of disasters, with the levels of the risk-free rates matching closely to the interest rates in the data (with the caveat that the empirical interest rates are nominal rather than real), and matching the spread between the rates closely at about 2% per annum. Consequently, the Sharpe ratio is roughly as high as in the data on average (around 0.4 on average in no disaster samples and just under 0.3 overall). However, the model does not completely rely on the peso-problem explanation of the carry trade profitability, as even in the samples including disasters the average carry trade return is essentially of the same magnitude. The volatility of exchange rates (and therefore currency carry strategy returns) in the model averaged over no disaster samples matches closely to the empirical volatility of the IMX returns for G10 currencies, at just over 7% per annum. This is below the unconditional mean and median over the simulated samples of 10.5% but slightly above the full sample median. Similarly, volatilities of consumption and output growth in the no disaster samples on average match those in the data, and are roughly between the means and the medians of the unconditional distributions. Thus, the model's ability to match unconditional currency risk premia does not rely on an unreasonably large magnitude (and probability) of a rare disaster.

The trade cost coefficients combined with the shipping sector dynamics imply that the fraction of total exports of the final good that is lost to transportation frictions is substantial, at close to 40% (but much smaller, around 11%, for commodities). These costs appear large but are in fact well within the range of values estimated by Anderson and van Wincoop (2004). The dynamics of the trade costs produced by the model are much less volatile than those observed in the data (we use the Baltic Dry Shipping index as our empirical proxy)[15].

Graveline, and Zviadadze (2012) points to the importance of crash risk in explaining jointly the carry trade risk premia and prices of currency options.

[15]The parameters governing mean reversion of the commodity production and shipping prices are chosen so

The calibrated model does feature predictability of carry trade returns with trade costs, as well as commodity prices, as described qualitatively in section 2.5.[16] We report average coefficients from predictive regressions analogous to those estimated in section 3.8, with standard errors constructed as standard deviations of point estimates across the simulated samples, in Table 14. As in the data, there is statistically significant relation between future currency returns and lagged changes in trade costs and commodity prices at short horizons (1- to 3-month returns), which fades away at longer horizons. Thus, the model appears to be able to rationalize the initially puzzling evidence of pro-cyclical predictability of carry trade strategies.

5 Conclusion

We present new evidence on the relation of the currency carry trade profits to the patterns in international trade: countries that specialize in exporting basic goods such as raw commodities tend to exhibit high interest rates where as countries primarily exporting finished goods have lower interest rates on average. These interest rate differences translate almost entirely into average returns on currency carry trade strategies. We propose a novel mechanism that helps rationalize these findings: convex shipping costs combined with time-varying capacity of the shipping industry. Nonlinearity of the shipping costs implies that the consumption - and therefore the SDF - of the country producing the consumption good is more sensitive to productivity shocks, and is thus riskier. Our model's empirical predictions are strongly supported in the data, while the quantitative analysis suggests that our mechanism may provide a fruitful direction for understanding the interaction between currency risk premia and the macroeconomy.

that the commodity production reverts more quickly than the shipping capital. This is broadly consistent with the behavior of commodity prices and shipping costs after the crisis, and also consistent with Bessembinder, Coughenor, Seguin, and Smoller (1995) who document relatively rapid mean reversion in commodity prices, and Kalouptsidi (2011) who emphasizes the long production lags in the shipping industry.

[16]The direct prediction of our model is that the *level* of trade costs comoves with the currency risk premium, and should therefore should forecast carry trade returns. We use changes in trade costs and commodity prices in order to avoid spurious predictability due to the small sample bias (e.g., Bekaert, Hodrick, and Marshall (1997), Stambaugh (1999)), since both variables are highly persistent, while their growth rates are only moderately autocorrelated.

Table 6: **Carry Trade Alphas and IMX**

	Trade-sorted Portfolios						
	1	2	3	4	5	6	IMX
β_j	-0.32**	-0.07	-0.00	0.12	0.26**	0.68**	
	(0.06)	(0.08)	(0.09)	(0.08)	(0.09)	(0.06)	
α_j	1.58	2.51	3.38	1.59	1.17	1.58	
	(1.51)	(1.86)	(2.01)	(1.51)	(1.82)	(1.51)	
R^2	0.14	0.01	0.00	0.02	0.07	0.42	1.00

	Portfolios sorted on Current Forward Discounts						
	1	2	3	4	5	6	HML_{FX}
β_j	-0.09	0.05	0.08	0.05	0.15	0.17	0.26**
	(0.06)	(0.06)	(0.07)	(0.08)	(0.09)	(0.11)	(0.08)
α_j	-2.02	-0.72	0.72	3.77*	1.81	6.31**	8.33**
	(1.43)	(1.35)	(1.45)	(1.61)	(1.84)	(2.26)	(1.92)
R^2	0.01	0.01	0.01	0.00	0.02	0.02	0.07

	Portfolios sorted on Past Mean Forward Discounts						
	1	2	3	4	5	6	$UHML_{FX}$
β_j	-0.12*	0.47**	0.44**	0.54**	0.87**	0.91**	1.03**
	(0.06)	(0.08)	(0.06)	(0.07)	(0.07)	(0.08)	(0.06)
α_j	-0.15	-2.52	-0.97	-2.11	-0.60	-0.43	-0.28
	(1.80)	(2.41)	(1.54)	(1.84)	(1.86)	(2.48)	(1.95)
R^2	0.02	0.15	0.26	0.28	0.49	0.39	0.55

This table reports regressions of the form

$$RX_t^j = \alpha_j + \beta_j IMX_t + \epsilon_t^j$$

where portfolios j are the six component portfolios of IMX (trade-based sort), HML_{FX} (conditional interest rate sort), and $UHML_{FX}$ (sort on an unconditional average forward discount over the period 1984 - 1995). Returns are monthly from 1988-2012 for the IMX and HML_{FX} sorts, and monthly $1995-2012$ for the $UHML_{FX}$ sorts. Standard errors are White (1980).

Table 7: Riskiness of aggregate consumption baskets, data

Portfolio	σ	β_{IMX}
Commodity Producers	0.92	0.013
	(0.09)	(0.009)
Final Goods Producers	1.40	0.033
	(0.18)	(0.014)

This table reports summary statistics from consumption portfolios formed on a country's commodity-making or final-good-producing status. The data are quarterly and taken from the OECD. The countries for which data for consumption and forward contracts are available are ranked according to the average import export measure used in constructing IMX. The commodity and final goods producers are the top and bottom third respectively. Consumption growth is calculated as the average growth rate of consumption weighted by the GDP of each country. Annualized standard deviations are estimated using quarterly growth rates for the time period from first-quarter 1988 until fourth-quarter 2012. Consumption betas are with respect to the quarterly IMX return. Standard errors are bootstrapped for the standard deviations and OLS for the IMX Betas.

Table 8: **Carry Trade Contemporaneous Relations**

Panel I: Import Ratio Sort (IMX)			
	(1)	(2)	(3)
	IMX	IMX	IMX
Δbdi_t	0.030**		0.019*
	(0.014)		(0.010)
ΔCRB_t		0.345**	0.323**
		(0.070)	(0.063)
Cons.	4.415**	4.039**	3.799**
	(1.836)	(1.737)	(1.710)
Obs	304	304	304
R^2	0.034	0.133	0.146

Panel II: Conditional Interest Rate Sort (HML_{FX})			
	(1)	(2)	(3)
	HML_{FX}	HML_{FX}	HML_{FX}
Δbdi_t	0.022*		0.017*
	(0.012)		(0.010)
ΔCRB_t		0.172*	0.152*
		(0.090)	(0.083)
Cons.	8.645**	8.564**	8.354**
	(1.904)	(1.908)	(1.908)
Obs	304	304	304
R^2	0.017	0.031	0.041

Panel III: HML_FX net of position in IMX ($CHML_{FX}$)			
	(1)	(2)	(3)
	$CHML_{FX}$	$CHML_{FX}$	$CHML_{FX}$
Δbdi_t	0.003		0.004
	(0.008)		(0.009)
ΔCRB_t		-0.010	-0.015
		(0.057)	(0.058)
Cons.	7.554**	7.628**	7.582**
	(1.788)	(1.804)	(1.809)
Obs	304	304	304
R^2	0.000	0.000	0.001

Regressions of currency carry-trade strategy returns on contemporaneous innovations in the Baltic Dry Index (BDI) and contemporaneous changes of the CRB All Commodity spot index. IMX, HML_{FX}, and $CHML_{FX}$ are as defined in Table 5. ΔCRB_t is the change in the long of the CRB index and Δbdi_t is the change in the log of the BDI. All data is monthly from 1/1988 to 12/2012. Standard errors are White (1980).

Table 9: **IMX and Commodity Prices**

Index	IMX Beta
CRB Textile Index	0.156
	(0.062)
Gold	0.314
	(0.145)
CRB Foodstuff Index	0.338
	(0.105)
CRB Livestock Index	0.376
	(0.156)
CRB Spot Commodity Index	0.386
	(0.125)
CRB Fats and Oils Index	0.406
	(0.198)
CRB Raw Industrials Index	0.413
	(0.149)
Platinum	0.558
	(0.174)
Silver	0.692
	(0.207)
CRB Industrial Metals Index	0.775
	(0.281)
Energy Goods	0.953
	(0.245)

This table reports β_i from regressions of the form

$$\frac{P_t^i}{P_{t-1}^i} = \alpha_i^c + \beta_i^c IMX_t + \varepsilon_t^i$$

where i are different commodity price indices and selected individual commodities. Seven of the indices are from the Commodity Research Board and represent changes in spot prices of different classes of commodities. In addition, an index of energy commodities is constructed using data from the CRB on the spot prices of Propane, Heating Oil, Natural Gas, and Crude Oil. Finally, percentage changes of three metals: platinum, silver, and gold, are included individually. Regressions are monthly from from 1988 - 2012. Standard errors are White (1980).

Table 10: **Predicting the Carry-Trade with the BDI**

Panel I: Import Ratio Sort (IMX)					
		(1)	(2)	(3)	(4)
		IMX	IMX	IMX	IMX
	Horizon:	1-month	3-month	6-month	12-month
$\Delta bdi_{t-4,t-1}$		0.152**	0.093**	0.007	-0.017
		(0.058)	(0.034)	(0.029)	(0.021)
Observations		304	302	299	293
R^2		0.041	0.016	0.006	0.004

Panel II: Conditional Interest Rate Sort (HML_{FX})					
		(1)	(2)	(3)	(4)
		HML_{FX}	HML_{FX}	HML_{FX}	HML_{FX}
	Horizon:	1-month	3-month	6-month	12-month
$\Delta bdi_{t-4,t-1}$		0.126*	0.077*	0.001	0.013
		(0.050)	(0.031)	(0.026)	(0.024)
Observations		304	302	299	293
R^2		0.027	0.016	0.001	0.000

Panel III: HML_{FX} net of position in IMX ($CHML_{FX}$)					
		(1)	(2)	(3)	(4)
		$CHML_{FX}$	$CHML_{FX}$	$CHML_{FX}$	$CHML_{FX}$
	Horizon:	1-month	3-month	6-month	12-month
$\Delta bdi_{t-4,t-1}$		0.036	0.022	-0.000	0.030
		(0.035)	(0.025)	(0.025)	(0.019)
Observations		304	302	299	293
R^2		0.003	0.004	0.000	0.003

Regressions of currency carry-trade strategy returns on the lag of the innovation to the BDI. IMX, HML_{FX}, and $CHML_{FX}$ are as defined in Table 5. $\Delta bdi_{t-4,t-1}$ is the change in the log of the BDI over the three months prior to the current period. All data is monthly. Standard errors in the parentheses are Newey-West with the number of lags equal to the horizon. For the 1 month horizon 3 lags are used.

Table 11: **Predicting the G10 Carry-Trade with Commodity Prices**

Panel I: Import Ratio Sort (IMX^{G10})				
	(1)	(2)	(3)	(4)
	IMX^{G10}	IMX^{G10}	IMX^{G10}	IMX^{G10}
Horizon:	1-month	3-month	6-month	12-month
$\Delta CRBIM_{t-4,t-1}$	0.518**	0.396*	0.058	-0.038
	(0.239)	(0.228)	(0.107)	(0.078)
Observations	304	302	299	293
R^2	0.015	0.025	0.001	0.001

Panel II: Conditional Interest Rate Sort (HML_{FX}^{G10})				
	(1)	(2)	(3)	(4)
	HML_{FX}^{G10}	HML_{FX}^{G10}	HML_{FX}^{G10}	HML_{FX}^{G10}
Horizon:	1-month	3-month	6-month	12-month
$\Delta CRBIM_{t-4,t-1}$	0.249	0.421**	0.042	-0.073
	(0.255)	(0.186)	(0.168)	(0.086)
Observations	304	302	299	293
R^2	0.003	0.021	0.000	0.002

Panel III: HML_{FX}^{G10} net of position in IMX^{G10} ($CHML_{FX}^{G10}$)				
	(1)	(2)	(3)	(4)
	$CHML_{FX}^{G10}$	$CHML_{FX}^{G10}$	$CHML_{FX}^{G10}$	$CHML_{FX}^{G10}$
Horizon:	1-month	3-month	6-month	12-month
$\Delta CRBIM_{t-4,t-1}$	-0.025	0.083	-0.007	-0.050
	(0.221)	(0.122)	(0.134)	(0.106)
Observations	304	302	299	293
R^2	0.000	0.001	0.000	0.001

Regressions of currency carry-trade strategy returns formed using the sample of G10 country currencies on the lagged growth rate in commodity prices. IMX^{G10}, HML_{FX}^{G10}, and $CHML_{FX}^{G10}$ are as defined in Table 5. $\Delta CRBIM_{t-4,t-1}$ is the change in the logarithm of the CRB Industrial Metals spot commodity price index over the three months prior to the current period. All data is monthly. Standard errors in the parentheses are Newey-West with the number of lags equal to the horizon. For the 1 month horizon 3 lags are used.

Table 12: **Calibration moments**

This table reports summary statistics generated by the model and compares them to data analogues from the G10 country set. The macroeconomic variables (consumption, output, exports) are time-aggregated quarterly. All of the financial variables (real interest rates, commodity prices, exchange rates, currency returns) are sampled monthly (monthly carry trade returns are based on continuously rolled-over positions in the model and one-month forward contract returns in the data). Real interest rates are calculated using 1 year lags of realized inflation to proxy for expected inflation. "AC" is the sample autocorrelation. The commodity country set includes Australia, Canada, New Zealand and Norway. The producer country set consists of Germany/Euro, Japan, Sweden, and Switzerland. All means and standard deviations are annualized, in percentage points. The model moments are averages across 10,000 simulated paths of 30 year length, reported as unconditional means and medians, as well as means conditional on "no disasters" - i.e., no jumps generating producer-country annual consumption declines greater than 5% over the 30-year period (disasters of such magnitude occur at least once in approximately 16% of simulated paths).

	Medians			**Means**			**Means, no disasters**			**Data**		
	Mean	Std	AC	Mean	Std	AC	Mean	Std	AC	Mean	Std	AC
Δy_{pt}	1.58	0.93	0.24	1.38	2.34	0.25	1.53	1.32	0.25	1.23	1.83	0.31
Δy_{ct}	1.56	0.70	0.60	1.34	1.72	0.58	1.50	0.98	0.56	2.84	0.97	0.43
Δc_{pt}	1.58	0.91	0.24	1.40	2.23	0.25	1.53	1.28	0.25	1.40	1.41	-0.20
Δc_{ct}	1.70	0.39	0.25	1.57	1.27	0.25	1.67	0.62	0.25	2.92	0.92	0.31
ΔX_t	1.58	0.94	0.24	1.37	2.43	0.25	1.52	1.35	0.25	3.21	10.21	0.02
r^f_{pt}	3.24	0.78	0.73	3.27	1.20	0.73	3.23	0.89	0.74	2.44	0.71	0.92
r^f_{ct}	6.64	0.28	0.82	6.49	0.58	0.82	6.60	0.36	0.83	4.65	0.58	0.94
$dRet_t$	2.74	6.95	0.09	2.37	10.56	0.09	2.70	7.49	0.09	2.86	7.62	0.04
dS_t	0.66	6.71	-0.01	0.89	10.62	-0.01	0.67	7.16	-0.01	-0.38	7.58	0.04
dP_t	-0.28	2.48	0.02	-0.39	5.19	0.03	-0.29	2.89	0.03	1.36	10.91	0.21
$d\tau_f(X_t, z_{kt})$	-0.54	3.01	0.23	-0.81	6.17	0.23	-0.57	3.46	0.24	6.15	56.25	0.33

Table 13: **Parameter values**

Parameter	**Value**	**Description**
λ	1	Relative Pareto weight
β	0.9	Cobb-Douglas producer-country labor share
γ	5	Relative risk aversion
ρ	0.001	Rate of time preference (annualized)
κ_0^c	0.01	Fixed commodity trade cost
κ_1^c	0.55	Variable commodity trade cost
κ_0^f	0.001	Fixed final trade cost
κ_1^f	0.75	Variable final trade cost
σ_p	0.0025	Productivity shock volatility (annualized)
σ_k	0.0001	Shipping shock volatility (annualized)
σ_c	0.0015	Commodity shock volatility (annualized)
μ	0.018	Uncompensated TFP growth rate (annualized)
ψ	0.01	Mean reversion of commodity supply (z_c to z_p)
ψ_k	0.00001	Mean reversion of shipping capacity (z_k to z_c)
η	1 per 25 years	jump frequency
α	1.1	Power tail of jump
Z_{min}	2%	Minimum jump size
Z_{max}	120%	Maximum jump size

Table 14: **Model predictive regressions**

This table reports regression statistics generated by the model. All regressions include an intercept (not reported) and are run on monthly simulated data (we report the mean across all simulations). All quantities and prices are in annualized units. The return horizon lengths denote cumulative horizon returns: $dRet(t, t+x)$, for example, denotes the x-month cumulative return. The regressors are three-month log differences in commodity prices and trade costs. Standard errors in the parentheses are estimated as standard deviations of point estimates across simulated samples.

	(1)	(2)	(3)	(4)
	$dRet_t$	$dRet_t$	$dRet_t$	$dRet_t$
Horizon:	1-month	3-month	6-month	12-month
$\Delta \log P_t$	1.88**	1.43*	0.99	0.55
	(0.87)	(0.78)	(0.66)	(0.56)
R^2	0.01	0.02	0.02	0.01
$\Delta\tau_f$	1.95**	1.49*	1.04	0.58
	(0.95)	(0.83)	(0.68)	(0.57)
R^2	0.01	0.02	0.02	0.01

References

Anderson, James E., and Eric van Wincoop, 2004, Trade costs, *Journal of Economic Literature* 42, 691–751.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters implied by equity index options, *Journal of Finance* 66, 1969–2012.

Backus, David, Patrick Kehoe, and Finn Kydland, 1992, International real business cycles, *Journal of Political Economy* 100, 745–75.

Bakshi, Gurdip, Peter Carr, and Liuren Wu, 2008, Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies, *Journal of Financial Economics* 87, 132–156.

Bakshi, Gurdip, and George Panayotov, 2012, Predictability of currency carry trades and asset pricing implications, *Journal of Financial Economics, forthcoming.*

Bakshi, Gurdip S, and Zhiwu Chen, 1997, Equilibrium valuation of foreign exchange claims, *Journal of Finance* 52, 799–826.

Bakshi, Gurdip S., George Panayotov, and Georgios Skoulakis, 2010, The baltic dry index as a predictor of global stock returns, commodity returns, and global economic activity, *SSRN eLibrary.*

Bansal, Ravi, and Ivan Shaliastovich, 2012, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies.*

Barro, Robert, 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.

Barro, Robert J., and Jose F. Ursua, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* 39, 255–350.

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, On biases in tests of the expectations hypothesis of the term structure of interest rates, *Journal of Financial Economics* 44(3), 30948.

Bessembinder, Hendrik, Jay F. Coughenor, Paul J. Seguin, and Margaret Monroe Smoller, 1995, Mean reversion in equilibrium asset prices: Evidence from the futures term structure, *The Journal of Finance* 50, 361–375.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, 2008, Carry trades and currency crashes, in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, ed.: *NBER Macroeconomics Annual* (University of Chicago Press).

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2008, Can peso problems explain the returns to the carry trade?, NBER Working Paper 14054.

Campbell, John Y., Karine Serfaty De Medeiros, and Luis M. Viceira, 2010, Global currency hedging, *Journal of Finance* 65, 87–121.

Chernov, Mikhail, Jeremy Graveline, and Irina Zviadadze, 2012, Crash risk in currency returns, working paper.

Colacito, R., and M. Croce, 2012, International asset pricing with recursive preferences, *Journal of Finance, forthcoming.*

Colacito, Riccardo, Mariano Massimiliano Croce, Steven Ho, and Philip Howard, 2013, BKK the EZ Way: An international production economy with recursive preferences, Working Paper, University of North Carolina.

Cole, Harold L., and Maurice Obstfeld, 1991, Commodity trade and international risk sharing : How much do financial markets matter?, *Journal of Monetary Economics* 28, 3–24.

Dumas, Bernard, 1992, Dynamic equilibrium and the real exchange rate in a spatially separated world, *Review of Financial Studies* 5, 153–80.

Eaton, Jonathan, Samuel Kortum, Brent Neiman, and John Romalis, 2011, Trade and the global recession, Working Paper 16666 National Bureau of Economic Research.

Farhi, Emmanuel, Samuel P. Fraiberger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan, 2009, Crash risk in currency markets, Working Paper.

Farhi, Emmanuel, and Xavier Gabaix, 2008, Rare disasters and exchange rates: A theory of the forward premium puzzle, Working Paper Harvard University.

Ferraro, Domenico, Barbara Rossi, and Kenneth S. Rogoff, 2011, Can Oil Prices Forecast Exchange Rates?, *Duke University Working Paper.*

Fitzgerald, Doireann, 2012, Trade costs, asset market frictions, and risk sharing, *American Economic Review* 102, 2700–2733.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *The Quarterly Journal of Economics* 127, 645–700.

Gourio, Francois, Michael Siemer, and Adrien Verdelhan, 2013, International risk cycles, *Journal of International Economics* 89, 471 – 484.

Greenwood, Robin, and Samuel Hanson, 2013, Waves in ship prices and investment, working paper, Harvard Business School.

Hassan, Tarek A., 2010, Country Size, Currency Unions, and International Asset Returns, *SSRN eLibrary.*

Hollifield, Burton, and Raman Uppal, 1997, An examination of uncovered interest rate parity in segmented international commodity markets, *Journal of Finance* 52, 2145–2170.

Jurek, Jakub W., 2009, Crash-Neutral Currency Carry Trades, Working paper.

Kalouptsidi, Myrto, 2011, Time to build and shipping prices, Yale University working paper.

Koijen, Ralph S.J., Lasse Heje Pedersen, Tobias J. Moskowitz, and Evert B. Vrugt, 2012, Carry, Chicago Booth working paper.

Lettau, Martin, Matteo Maggiori, and Michael Weber, 2013, Conditional risk premia in currency markets and other asset classes, Working Paper 18844 National Bureau of Economic Research.

Lewis, Karen K., 1995, Puzzles in international financial markets, in G. Grossman, and K. Rogoff, ed.: *Handbook of International Economics, Vol.III* (Elsevier: Amsterdam).

Longstaff, Francis A., and Monika Piazzesi, 2004, Corporate earnings and the equity premium, *Journal of Financial Economics* 74, 401–421.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency returns, *Review of Financial Studies* 24, 3731–3777.

——— , 2013, Countercyclical currency risk premia, *Journal of Financial Economics, forthcoming.*

Lustig, Hanno, and Adrien Verdelhan, 2007, The cross-section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.

Lyons, Richard K., 2001, *The Microstructure Approach to Exchange Rates* (MIT Press).

Martin, Ian, 2011, The Forward Premium Puzzle in a Two-Country World, *NBER Working Paper no. 17564.*

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012, Carry trades and global foreign exchange rate volatility, *Journal of Finance* 67 (2), 681–718.

Obstfeld, Maurice, and Kenneth Rogoff, 2001, The six major puzzles in international macroeconomics: is there a common cause?, in *NBER Macroeconomics Annual 2000, Volume 15* . pp. 339–412 (MIT press).

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Stathopoulos, Andreas, Andrea Vedolin, and Philippe Mueller, 2012, International correlation risk, working paper.

Verdelhan, Adrien, 2010, A habit-based explanation of the exchange rate risk premium, *Journal of Finance* 65, 123–145.

Wachter, Jessica, 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *The Journal of Finance* 68, 987–1035.

White, Halbert, 1980, A heteroskedasticity-consistent covariance estimator and direct test for heteroskedasticity, *Econometrica* 48, 817–838.

Appendix

5.1 Output

Commodity output y_{ct} equals the level of z_{ct}, so that the final good output dynamics are given by

$$
\begin{aligned}
y_{pt} &= z_{pt}[z_{ct}(1-\tau_c(z_{ct}, z_{kt}))]^{1-\beta} \\
&= z_{pt} I(z_{ct}, z_{kt})^{1-\beta} \\
dy_{pt} &= dz^c_{pt} I_t^{1-\beta} \\
&+ z_{pt}(1-\beta) I_t^{-\beta} I_c dz^c_{ct} + \frac{1}{2} z_{pt}(1-\beta)\left(I_t^{-\beta} I_{cc} - \beta I_t^{-\beta-1} I_c^2\right) dz^c_{ct} dz^c_{ct} \\
&+ z_{pt}(1-\beta) I_t^{-\beta} I_k dz^c_{kt} + \frac{1}{2} z_{pt}(1-\beta)\left(I_t^{-\beta} I_{kk} - \beta I_t^{-\beta-1} I_k^2\right) dz^c_{kt} dz^c_{kt} \\
&+ d\left(\sum_{0<s\leq t} (y_{ps} - y_{ps^-})\right) \\
&= z_{pt}\mu_p I_t^{1-\beta} dt + z_{pt}\sigma_p I_t^{1-\beta} dB_{pt} \\
&+ \left(z_{pt}(1-\beta) I_t^{-\beta} I_c z_{ct}\mu_{ct} + \frac{1}{2} z_{pt}(1-\beta)\left(I_t^{-\beta} I_{cc} - \beta I_t^{-\beta-1} I_c^2\right) z_{ct}^2 \sigma_c^2\right) dt \\
&+ \left(z_{pt}(1-\beta) I_t^{-\beta} I_k z_{kt}\mu_{kt} + \frac{1}{2} z_{pt}(1-\beta)\left(I_t^{-\beta} I_{kk} - \beta I_t^{-\beta-1} I_k^2\right) z_{kt}^2 \sigma_k^2\right) dt \\
&+ z_{pt}(1-\beta) I_t^{-\beta} I_c z_{ct}\sigma_c dB_{ct} + z_{pt}(1-\beta) I_t^{-\beta} I_k z_{kt}\sigma_k dB_{kt} + z_{pt} I_t^{1-\beta}(e^{Z_{N(t)}} - 1) dN_t \\
\Rightarrow \frac{dy_{pt}}{y_{pt^-}} &= \mu_p dt + \sigma_p dB_{pt} \\
&+ (1-\beta)\left[\frac{I_c}{I_t} z_{ct}\mu_{ct} + \frac{1}{2}\left(\frac{I_{cc}}{I_t} - \beta\frac{I_c^2}{I_t^2}\right) z_{ct}^2\sigma_c^2\right] dt + (1-\beta)\frac{I_c}{I_t} z_{ct}\sigma_c dB_{ct} \\
&+ (1-\beta)\left[\frac{I_k}{I_t} z_{kt}\mu_{kt} + \frac{1}{2}\left(\frac{I_{kk}}{I_t} - \beta\frac{I_k^2}{I_t^2}\right) z_{kt}^2\sigma_k^2\right] dt + (1-\beta)\frac{I_k}{I_t} z_{kt}\sigma_k dB_{kt} \\
&+ \left(e^{Z_{N(t)}} - 1\right) dN_t \\
&\mathrel{\hat{=}} \mu_{yt} dt + \sigma_{yt}^T dB_t + (e^{Z_{N(t)}} - 1) dN_t,
\end{aligned}
$$

where $I(z_{ct}, z_{kt})$ and its derivatives are defined as follows:

$$
\begin{aligned}
I_t = I(z_{ct}, z_{kt}) &= z_{ct}(1 - \tau_c(z_{ct}, z_{kt})) \\
I_c &= (1 - \kappa_0^c) - 2\kappa_1^c \frac{z_{ct}}{z_{kt}} \\
I_{cc} &= -2\kappa_1^c / z_{kt} \\
I_k &= \kappa_1^c \frac{z_{ct}^2}{z_{kt}^2} \\
I_{kk} &= -2\kappa_1^c \frac{z_{ct}^2}{z_{kt}^3}
\end{aligned}
$$

Commodity price dynamics are given by

$$
\begin{aligned}
P_t &= (1 - \beta) z_{pt} \left[z_{ct}(1 - \tau_c(z_{ct}, z_{kt})) \right]^{-\beta} \\
&= \frac{(1 - \beta) y_{pt}}{(1 - \tau_c(z_{ct}, z_{kt})) z_{ct}}
\end{aligned}
$$

5.2 Exports of final consumption good

Since in the general case the export function must be found numerically, it is convenient to restate equation (1) as

$$
\left[\xi_t(1 - \kappa_0^f - \kappa_1^f \xi_t)\right]^{-\gamma} \left(1 - \kappa_0^f - 2\kappa_1^f \xi_t\right) - \lambda \left[\exp\left(q_t + q_{kt}\right)\left(1 - \kappa_0^c - \kappa_1^c \exp\left(q_{kt}\right)\right)^{1-\beta} - \xi_t\right]^{-\gamma} = 0
$$

where $\xi_t = \frac{X_t}{z_{kt}} \triangleq \xi\left(q_t, q_{kt}\right)$ is exports of final good per unit of shipping capacity as a function of the two stationary state variables. Then the numerical solution for ξ_t can be interpolated for use in simulations.

In the special case of log utility ($\gamma = 1$) equation (1) simplifies to

$$
\kappa_1^f (2 + \lambda) X_t^2 - [z_{kt}(1 - \kappa_0^f)(1 + \lambda) + 2\kappa_1^f y_{pt}] X_t + (1 - \kappa_0^f) y_{pt} z_{kt} = 0.
$$

Solving this equation yields

$$X_t = \frac{z_{kt}(1-\kappa_0^f)(1+\lambda) + 2\kappa_1^f y_{pt} - \sqrt{[z_{kt}(1-\kappa_0^f)(1+\lambda) + 2\kappa_1^f y_{pt}]^2 - 4(1-\kappa_0^f) y_{pt} z_{kt} \kappa_1^f (2+\lambda)}}{2\kappa_1(2+\lambda)}$$

which is the only root that allows positive producer-country consumption. We can write

$$X_t = \frac{h(z_{ct}, z_{pt}, z_{kt}) - \sqrt{g(z_{ct}, z_{pt}, z_{kt})}}{2\kappa_1(2+\lambda)},$$

where

$$\begin{aligned} h(z_{ct}, z_{pt}, z_{kt}) &= z_{kt}(1-\kappa_0)(1+\lambda) + 2\kappa_1 z_{pt} I_t^{1-\beta}, \\ g(z_{ct}, z_{pt}, z_{kt}) &= h(z_{ct}, z_{pt}, z_{kt})^2 - 4(1-\kappa_0)\kappa_1(2+\lambda) z_{pt} I_t^{1-\beta} z_{kt}. \end{aligned}$$

The derivatives of the export function and its components follow:

$$X_i = \frac{h_i - \frac{1}{2} g^{-1/2} g_i}{2\kappa_1(2+\lambda)}, \ \forall\, i = \{c, p, k\}$$

$$X_{ii} = \frac{h_{ii} + \frac{1}{4} g^{-3/2} g_i^2 - \frac{1}{2} g^{-1/2} g_{ii}}{2\kappa_1(2+\lambda)}.$$

In the general CRRA case the derivatives of the export function can be found by implicit differentiation:

$$\frac{dX}{dz_i} = -\frac{g_{z_i}}{g_X} \ \text{ for } i \in c, p, k$$

$$\frac{d^2 X}{(dz_i)^2} = -\left(\frac{g_X \left(g_{z_i,X} \frac{dX}{dz_i} + g_{z_i,z_i} \right) - g_{z_i} \left(g_{X,X} \frac{dX}{dz_i} + g_{X,z_i} \right)}{(g_X)^2} \right)$$

By normalizing each partial differential by X_t and by Ito's lemma,

$$\begin{aligned}
dX_t(z_{ct}, z_{pt}, z_{kt}) &= X_{ct}X_t dz^c_{ct} + X_{pt}X_t dz^c_{pt} + X_{kt}X_t dz^c_{kt} \\
&\quad + \frac{1}{2}X_{cct}X_t dz^c_{ct}dz^c_{ct} + \frac{1}{2}X_{ppt}X_t dz^c_{pt}dz^c_{pt} + \frac{1}{2}X_{kkt}X_t dz^c_{kt}dz^c_{kt} \\
&\quad + d\left(\sum_{0<s\le t}(X_s - X_{s^-})\right) \\
\Rightarrow \frac{dX_t}{X_{t^-}} &= \left\{X_c\mu_{ct}z_{ct} + X_{pt}\mu_p z_{pt} + X_{kt}\mu_{kt}z_{kt} + \frac{1}{2}X_{cct}\sigma_c^2 z_{ct}^2 + \frac{1}{2}X_{ppt}\sigma_p^2 z_{pt}^2 + \frac{1}{2}X_{kkt}\sigma_k^2 z_{kt}^2\right\}dt \\
&\quad + X_{ct}\sigma_c z_{ct}dB_{ct} + X_{pt}\sigma_p z_{pt}dB_{pt} + X_{kt}\sigma_k z_{kt}dB_{kt} + d\left(\sum_{0<s\le t}\frac{X_s - X_{s^-}}{X_{t^-}}\right) \\
&\mathrel{\hat{=}} \mu_{Xt}dt + \sigma_{Xt}^T dB_t + \left(e^{\mathcal{J}_X} - 1\right)dN_t,
\end{aligned}$$

where $\mathcal{J}_X = \log\left(\xi(q_{t^-} + Z_{N(t)}, q_{kt^-})\right) - \log\left(\xi(q_{t^-}, q_{kt^-})\right)$, the log change in final goods exported.

5.3 Consumption

For the consumption allocations we have

$$\begin{aligned}
c_{pt} &= y_{pt} - X_t \\
\Rightarrow dc_{pt} &= dy^c_{pt} - dX^c_t + d\left(\sum_{0<s\le t}(c_{ps} - c_{ps^-})\right) \\
\Rightarrow \frac{dc_{pt}}{c_{pt^-}} &= \frac{1}{c_{pt^-}}\left(\mu_{yt} - \mu_{Xt}\right)dt + \frac{1}{c_{pt^-}}\left(\sigma_{yt}^T - \sigma_{Xt}^T\right)dB_t + d\left(\sum_{0<s\le t}\frac{c_{ps} - c_{ps^-}}{c_{pt^-}}\right) \\
&\mathrel{\hat{=}} \mu_{cpt}dt + \sigma_{cpt}^T dB_t + \left(e^{\mathcal{J}_p} - 1\right)dN_t
\end{aligned}$$

for the final good producer, and

$$
\begin{aligned}
c_{ct} &= X_t \left(1 - \kappa_0^f - \kappa_1^f \frac{X_t}{z_{kt}}\right) \\
dc_{ct} &= (1-\kappa_0^f) dX_t^c - \kappa_1^f d\left(\frac{(X_t^2)^c}{z_{kt}^c}\right) + d\left(\sum_{0<s\le t} (c_{cs} - c_{cs^-})\right) \\
\Rightarrow \frac{dc_{ct}}{c_{ct^-}} &= \frac{1}{c_{ct^-}} \left\{ \mu_{Xt}(1-\kappa_0^f) - \kappa_1^f \left[\frac{1}{z_{kt}}(2X_t\mu_{Xt} + \sigma_{Xt}^T\sigma_{Xt}) - \frac{X_t^2}{z_{kt}}(\mu_{kt} - \sigma_k^2) - 2X_t X_{kt}\sigma_k^2\right]\right\} dt \\
&\quad + \frac{1}{c_{ct^-}}(1-\kappa_0^f)\sigma_{Xt}^T dB_t - \frac{1}{c_{ct^-}}\kappa_1^f \frac{2X_t}{z_{kt}}\sigma_{Xt}^T dB_t - \frac{1}{c_{ct^-}}\kappa_1^f \frac{X_t^2}{z_{kt}}\sigma_k dB_{kt} \\
&\quad + d\left(\sum_{0<s\le t} \frac{c_{cs} - c_{cs^-}}{c_{ct^-}}\right) \\
&\quad \mathrel{\mathring{=}} \mu_{cct} dt + \sigma_{cct}^T dB_t + \left(e^{\mathcal{J}_c} - 1\right) dN_t
\end{aligned}
$$

for the commodity producer.

5.4 Risk-free rates

In order to compute risk-free rates the expected growth rate of marginal utility conditional on a jump occurring must be computed as a function of the state variables. Let

$$
\begin{aligned}
\mathbb{E}_Z\left[e^{-\gamma \mathcal{J}_c}\right] &= \mathbb{E}_Z\left(\frac{\xi(q_{t^-} + Z, q_{kt^-})(1 - \kappa_0^f - \kappa_1^f \xi(q_{t^-} + Z, q_{kt^-}))}{\xi(q_{t^-}, q_{kt^-})(1 - \kappa_0^f - \kappa_1^f \xi(q_{t^-}, q_{kt^-}))}\right)^{-\gamma} \\
&\mathrel{\mathring{=}} \zeta_c(q_{t^-}, q_{kt^-}),
\end{aligned}
$$

since the distribution of jump sizes is time invariant. Similarly, let

$$
\begin{aligned}
\mathbb{E}_Z\left[e^{-\gamma \mathcal{J}_p}\right] &= \mathbb{E}_Z\left(\frac{\exp(q_{t^-} + Z + q_{kt^-})(1 - \kappa_0^c - \kappa_1^c \exp(q_{kt^-}))^{1-\beta} - \xi(q_{t^-} + Z, q_{kt^-})}{\exp(q_{t^-} + q_{kt^-})(1 - \kappa_0^c - \kappa_1^c \exp(q_{kt^-}))^{1-\beta} - \xi(q_{t^-}, q_{kt^-})}\right)^{-\gamma} \\
&\mathrel{\mathring{=}} \zeta_p(q_{t^-}, q_{kt^-}).
\end{aligned}
$$

These functions can be evaluated by integrating over the distribution of jump sizes Z given by the pdf $\varphi(Z) = \frac{\alpha Z_{min}^{\alpha} x^{-\alpha-1}}{1-\left(\frac{Z_{min}}{Z_{max}}\right)^{\alpha}}$; this is done numerically using Gaussian quadrature.

5.5 Exchange rate

Since the spot exchange rate is defined as

$$S_t = \lambda \left(\frac{c_{pt}}{c_{ct}} \right)^{-\gamma} = \left(1 - \kappa_0^f - 2\kappa_1^f \frac{X_t}{z_{kt}} \right),$$

we can derive the dynamic evolution of exchange rate changes as

$$\begin{aligned}
dS_t &= -2\kappa_1^f \left[\left(\frac{1}{z_{kt}} dX_t - \frac{X_t}{z_{kt}} (\mu_{kt} - \sigma_k^2) dt - \frac{X_t}{z_{kt}} \sigma_k dB_{kt} - X_{kt}\sigma_k^2 dt \right) + \frac{1}{z_{kt}} d \left(\sum_{0<s\leq t} (X_s - X_{s^-}) \right) \right] \\
&= -2\kappa_1^f \left[\frac{X_t}{z_{kt}} \left\{ X_{ct}\mu_{ct} z_{ct} + X_{pt}\mu_p z_{pt} + X_{kt}\mu_{kt} z_{kt} + \frac{1}{2} X_{cct}\sigma_c^2 z_{ct}^2 + \frac{1}{2} X_{ppt}\sigma_p^2 z_{pt}^2 + \frac{1}{2} X_{kkt}\sigma_k^2 z_{kt}^2 \right\} dt \right. \\
&\quad - \frac{X_t}{z_{kt}} (\mu_{kt} - \sigma_k^2) dt - X_t X_{kt} \sigma_k^2 dt + \frac{X_t}{z_{kt}} X_{ct}\sigma_c z_{ct} dB_{ct} + \frac{X_t}{z_{kt}} X_{pt}\sigma_p z_{pt} dB_{pt} + \frac{X_t}{z_{kt}} (X_{kt} - 1) \sigma_k dB_{kt} \\
&\quad \left. + \left(\xi \left(q_{t^-} + Z_{N(t)}, q_{kt^-} \right) - \xi \left(q_{t^-}, q_{kt^-} \right) \right) dN(t) \right] \\
\Rightarrow \frac{dS_t}{S_{t^-}} &\mathrel{\hat{=}} \mu_{St} dt + \sigma_{St}^T dB_t + \left(e^{\mathcal{J}_S} - 1 \right) dN(t),
\end{aligned}$$

where $\mathcal{J}_S = \log \left(1 - \kappa_0^f - 2\kappa_1^f \xi \left(q_{t^-} + Z_{N(t)}, q_{kt^-} \right) \right) - \log \left(1 - \kappa_0^f - 2\kappa_1^f \xi \left(q_{t^-}, q_{kt^-} \right) \right)$

5.6 Expected Returns

Let

$$E \left[dRet_t | \mathcal{F}_t \right] = E \left[\frac{dS_t}{S_{t^-}} \frac{d\pi_{pt}}{\pi_{pt^-}} | \mathcal{F}_t \right] \mathrel{\hat{=}} \mu_t^{FX} dt,$$

where μ_t^{FX} is the instantaneous conditional currency risk premium, which can be calculated as

$$\mu_t^{FX} = -\gamma \sigma_{St}^T \sigma_{cpt} + \eta \mathbb{E}_Z \left[\left(e^{\mathcal{J}_S} - 1 \right) \left(e^{-\gamma \mathcal{J}_p} - 1 \right) \right].$$

Figure 8 displays the final good trade costs τ_f and the conditional currency risk premium μ^{FX} as functions of the two cointegrating residuals q_t and q_t^k, evaluated at $q_t = 0$, so that a higher q_t^k due to large output of the final good relative to the available shipping capacity translates into high shipping costs and high expected excess returns.

Figure 8: Trade Costs and Currency Risk Premium